arXiv:1406.6169v1 [cs.DS] 24 Jun 2014

Fault Tolerant Approximate BFS Structures

Merav Parter *† David Peleg * ‡

June 25, 2014

Abstract

A *fault-tolerant* structure for a network is required to continue functioning following the failure of some of the network's edges or vertices. This paper addresses the problem of designing a *fault-tolerant* (α, β) approximate BFS structure (or `FT-ABFS` *structure* for short), namely, a subgraph H of the network G such that subsequent to the failure of some subset F of edges or vertices, the surviving part of H still contains an *approximate* BFS spanning tree for (the surviving part of) G, satisfying $\text{dist}(s, v, H \setminus F) \leq \alpha \cdot \text{dist}(s, v, G \setminus F) + \beta$ for every $v \in V$.

We first consider *multiplicative* $(\alpha, 0)$ `FT-ABFS` structures resilient to a failure of a single edge and present an algorithm that given an n-vertex unweighted undirected graph G and a source s constructs a $(3, 0)$ `FT-ABFS` structure rooted at s with at most $4n$ edges (improving by an $O(\log n)$ factor on the near-tight result of [3] for the special case of edge failures). Assuming at most f edge failures, for constant integer $f > 1$, we prove that there exists a (poly-time constructible) $(3(f+1), (f+1)\log n)$ `FT-ABFS` structure with $O(fn)$ edges.

We then consider *additive* $(1, \beta)$ `FT-ABFS` structures. In contrast to the linear size of $(\alpha, 0)$ `FT-ABFS` structures, we show that for every $\beta \in [1, O(\log n)]$ there exists an n-vertex graph G with a source s for which any $(1, \beta)$ `FT-ABFS` structure rooted at s has $\Omega(n^{1+\epsilon(\beta)})$ edges, for some function $\epsilon(\beta) \in (0, 1)$. In particular, $(1, 3)$ `FT-ABFS` structures admit a lower bound of $\Omega(n^{5/4})$ edges. These lower bounds demonstrate an interesting dichotomy between multiplicative and additive

*Department of Computer Science and Applied Mathematics. The Weizmann Institute of Science, Rehovot, Israel. E-mail: {merav.parter,david.peleg}@ weizmann.ac.il. Supported in part by the Israel Science Foundation (grant 894/09), the United States-Israel Binational Science Foundation (grant 2008348), the I-CORE program of the Israel PBC and ISF (grant 4/11), the Israel Ministry of Science and Technology (infrastructures grant), and the Citi Foundation.

†Recipient of the Google European Fellowship in distributed computing; research is supported in part by this Fellowship.

‡An extended abstract of this paper has appeared in the proceedings of the 2014 ACM-SIAM Symposium on Discrete Algorithms.

spanners; whereas $(\alpha, 0)$ FT-ABFS structures of size $O(n)$ exist (for $\alpha \geq 3$), their additive counterparts, $(1, \beta)$ FT-ABFS structures, are of super-linear size. Our lower bounds are complemented by an upper bound, showing that there exists a poly-time algorithm that for every n-vertex unweighted undirected graph G and source s constructs a $(1, 4)$ FT-ABFS structure rooted at s with at most $O(n^{4/3})$ edges.

1 Introduction

Background and Motivation. Fault-tolerant subgraphs are subgraphs designed to maintain a certain desirable property in the presence of edge or vertex failures. This paper focuses on the property of containing a BFS tree with respect to some source s. A *fault tolerant BFS structure* (or FT-BFS *structure*) resistant to a single edge failure is a subgraph $H \subseteq G$ satisfying that $\text{dist}(s, v, H \setminus \{e\}) = \text{dist}(s, v, G \setminus \{e\})$ for every vertex $v \in V$ and edge $e \in E$.

To motivate our interest in such structures, consider a situation where it is required to lease a subnetwork of a given network, which will provide short routes from a source s to all other vertices. In a failure-free environment one can simply lease a BFS tree rooted at s. However, if links might disconnect, then one must prepare by leasing a larger set of links, and specifically an FT-BFS structure. Moreover, taking costs into account, this example also motivates our interest in constructing *sparse* FT-BFS structure.

This question has recently been studied by us in [15]. Formally, a spanning graph $H \subseteq G$ is an *f edge (resp., vertex) fault-tolerant BFS (*FT-BFS*) structure* for G with respect to the source $s \in V$ iff for every $v \in V$ and every set $F \subseteq E(G)$ (resp., $F \subseteq V$), $|F| \leq f$, it holds that $\text{dist}(s, v, H \setminus F) = \text{dist}(s, v, G \setminus F)$. It is shown in [15] that for every graph G and source s there exists a (poly-time constructible) 1-edge FT-BFS structure H with $O(n^{3/2})$ edges. This result is complemented by a matching lower bound showing that for every sufficiently large integer n, there exist an n-vertex graph G and a source $s \in V$, for which every 1-edge FT-BFS structure is of size $\Omega(n^{3/2})$. Hence *exact* FT-BFS structures may be rather expensive.

This last observation motivates the approach of resorting to *approximate* distances, in order to allow the design of a sparse subgraph with properties resembling those of an FT-BFS structure. The current paper aims at exploring this approach, focusing mainly on subgraphs that contain *approximate* BFS structures and are resistant to a single edge failure. Formally, given an unweighted undirected n-vertex graph $G = (V, E)$ and a source $s \in V$, the subgraph $H \subseteq G$ is an f-edge (resp., vertex) (α, β) FT-ABFS structure with respect to s if for every vertex $v \in V$ and every set $F \subseteq E(G)$ (resp., $F \subseteq V$),

$|F| \leq f$,

$$\mathrm{dist}(s, v, H \setminus F) \leq \alpha \cdot \mathrm{dist}(s, v, G \setminus F) + \beta \ .$$

(An (α, β) FT-ABFS structure is a fault-tolerant BFS (FT-BFS) structure if $\alpha = 1$ and $\beta = 0$.) We show that this relaxed requirement allows structures that are sparser than their exact counterparts.

Approximate BFS tree structures can also be compared against a different type of structures, namely, fault-tolerant *spanners*. Given an n-vertex graph $G = (V, E)$, the subgraph $H \subseteq G$ is an *f-edge fault-tolerant (α, β) spanner* of G if for every two vertices $v, w \in V$ and every set $F \subseteq E(G)$, $|F| \leq f$, we have $\mathrm{dist}(v, w, H \setminus F) \leq \alpha \cdot \mathrm{dist}(v, w, G \setminus F) + \beta$. Observe that the union of (α, β) FT-ABFS structures with respect to every source $s \in V$ forms an (all-pairs) fault tolerant (α, β) spanner for G. In fact, FT-ABFS structures can be viewed as *single source spanners*. Algorithms for constructing an f-vertex fault tolerant $(2k-1)$ spanner of size $O(f^2 k^{f+1} \cdot n^{1+1/k} \log^{1-1/k} n)$ and an f-edge fault tolerant $2k-1$ spanner of size $O(f \cdot n^{1+1/k})$ for a given n-vertex graph G were presented in [8]. A randomized construction attaining an improved tradeoff for vertex fault-tolerant spanners was then presented in [11].

For the case of f edge failures for constant $f \geq 1$, we show (in Sec. 2) that there exists a poly-time algorithm that for every n-vertex graph constructs a $(3(f+1), (f+1)\log n)$ FT-ABFS structure H with $O(fn)$ edges overcoming up to f edge faults. For the special case of a single edge failure ($f = 1$), we get a somewhat stronger result, namely, that for every n-vertex graph G and source s, there is a (poly-time constructible) $(3, 0)$ FT-ABFS structure with at most $4n$ edges, thus improving on the near-tight construction of [3] by a $O(\log n)$ factor for the special case of $\alpha = 3$ and edge failures.

This result is to be contrasted with two different structures: the (single-source) fault tolerant exact FT-BFS structure of [15], and the (all-pairs) fault tolerant $(3, 0)$ spanner of [8], which both contain $\Theta(n^{3/2})$ edges. This implies that using FT-ABFS structures is more efficient than using fault-tolerant spanners even if it is necessary to handle not a single source s but a set $S \subseteq V$ of sources where $|S| = \Omega(n^\epsilon)$ for $\epsilon < 1/2$; a collection of approximate (α, β) FT-ABFS structures rooted at each of the sources $s \in S$ will still be cheaper than a fault-tolerant spanner.

Additive fault tolerant $(1, \beta)$ spanners were recently defined and studied by [6], establishing the following general result. For a given n-vertex graph G, let H_1 be an ordinary additive $(1, \beta)$ spanner for G and H_2 be a fault tolerant $(\alpha, 0)$ spanner for G resilient against up to f edge faults. Then $H = H_1 \cup H_2$ is a $(1, \beta(f))$ additive fault tolerant spanner for G (for up to f edge faults) for $\beta(f) = O(f(\alpha + \beta))$. In particular, fixing the number of H edges to be $O(n^{4/3})$ and the number of faults to $f = 1$ yields an additive stretch of 38 (See [6]; Cor. 1).

When considering FT-BFS structures with an *additive* stretch, namely, $(1, \beta)$ FT-ABFS

structures, the improvement is less dramatic compared to the size of the single-source exact or the all-pairs approximate variants. In Sec. 3, we show that for every additive stretch $\beta \in [1, \log n]$, there exists a superlinear lower bound on the size of the FT-ABFS structure with additive stretch β, i.e., $\Omega(n^{1+\epsilon(\beta)})$. These new lower bound constructions are independent of the correctness of Erdös conjecture. Importantly, our results reveal an interesting dichotomy between *multiplicative* $(\alpha, 0)$ FT-ABFS and *additive* $(1, \beta)$ FT-ABFS structures: whereas every graph G contains a (poly-time constructible) $(3, 0)$ FT-ABFS structure rooted at $s \in V$ of size $\Theta(n)$, there exist an n-vertex graph G and a source $s \in V$ for which every $(1, \beta)$ FT-ABFS structure contains a *super-linear* number of edges. For example, for additive stretch $\beta = 3$, we have a lower bound construction with $\Omega(n^{5/4})$ edges.

On the positive side, in Sec. 4 we complement those results by presenting a (rather involved) poly-time algorithm that for any given n-vertex graph G and source s constructs a $(1, 4)$ FT-ABFS structure with $O(n^{4/3})$ edges (hence improving the additive stretch of the (all-pairs) fault tolerant additive spanner with $O(n^{4/3})$ edges of [6] from 38 to 4). This algorithm is inspired by the (non-fault-tolerant) additive spanner constructions of [4, 9, 10]. The main technical contribution of our algorithm is in adapting the path-buying strategy used therein to failure-prone settings. So far, the correctness and size analysis of this strategy heavily relied on having a fault-free input graph G. We show that by a proper construction of the sourcewise replacement paths, the path-buying technique can be extended to support the construction even in the presence of failures.

Related work. FT-BFS structures are closely related to the notion of *replacement paths*. For a source s, a target vertex v and an edge $e \in G$, a *replacement path* is the shortest $s - v$ path $P_{s,v,e}$ that does not go through e. An FT-BFS structure is composed of a collection consisting of a replacement path $P_{s,v,e}$ for every target $v \in V$ and edge $e \in E$. Analogously, the notion of FT-ABFS structures is closely related to the problem of constructing *approximate replacement paths* [2, 7, 5], and in particular to its *single source* variant studied in [3]. That problem requires to compute a collection $\mathcal{P}_s$ consisting of an approximate $s - t$ replacement path $P_{s,t,e}$ for every $t \in V$ and every failed edge e that appears on the $s - t$ shortest-path in G, such that $|P_{s,t,e}| \leq \alpha \cdot \text{dist}(s, t, G \setminus \{e\})$. In the resulting *fault tolerant distance oracle*, in response to a query (s, t, F) consisting of an $s - t$ pair and a set F of failed edges or vertices (or both), the oracle $\mathcal{S}$ must return the distance between s and t in $G' = G \setminus F$. Such a structure is sometimes called an *F-sensitivity distance oracle*. The focus is on both fast preprocessing time, fast query time and low space. An approximate *single source fault tolerant distance oracle* has been first studied at [3], which proposed an $O(n \log n/\epsilon^3)$ space data structure that can report a $(1 + \epsilon)$ approximate shortest path for any $\epsilon > 0$. An additional by-

product of the data structure of [3] is the construction of an $(1+\epsilon, 0)$ FT-ABFS structure with $O(n/\epsilon^3 + n \log n)$ edges. Setting $\epsilon = 2$, this yields a $(3, 0)$ FT-ABFS structure with $O(n \log n)$ edges. Hence our $(3, 0)$ FT-ABFS structure construction with at most $3n$ edges improves that construction by a factor of $O(\log n)$ for the case of single edge failure (the construction of [3] supports the case of vertex failures as well).

It is important to note that the literature on approximate replacement paths (cf. [2, 5]) mainly focuses on *time-efficient* computation of the these paths, as well as their efficient maintenance within distance oracles. In contrast, the main concern in the current paper is with optimizing the *size* of the resulting fault tolerant structure that contains the collection of approximate replacement paths.

Moreover, this paper considers both multiplicative and additive stretch, whereas the long line of existing approximate distance oracles concerned mostly multiplicative (and not additive) stretch, with the exception of [16]. To illustrate the dichotomy between the additive and multiplicative setting, consider the issue of lower bounds for additive FT-ABFS structures. In the all-pairs fault-free setting, the best known lower bound for additive spanners is based on the *girth conjecture* of Erdös [12], stating that there exist n-vertex graphs with $\Omega(n^{1+1/k})$ edges and girth (minimum cycle length) $2k+2$ for any integer k. Removing any edge in such a graph increases the distance between its endpoints from 1 to $2k+1$, hence any $(1, \beta)$ spanner with $\beta \leq 2k-1$ must have $\Omega(n^{1+1/k})$ edges. This conjecture is settled only for $k = 1, 2, 3, 5$ (see [18]). In [19], Woodruff presented a lower bound for additive spanners matching the girth conjecture bounds but independent of the correctness of the conjecture. More precisely, he showed the existence of graphs for which any spanner of size $O(k^{-1}n^{1+1/k})$ has an additive stretch of at least $2k-1$, hence establishing a lower bound of $\Omega(k^{-1}n^{1+1/k})$ on the size of additive spanners. The lower bound constructions of [19] are formed by appropriately gluing together certain complete bipartite graphs. Since for every n-vertex graph G there exists a (poly-time constructible) multiplicative spanner of size $O(n^{1+1/k})$ and stretch $\alpha = 2k-1$, so far there has been no theoretical indication for a dichotomy between additive and multiplicative spanners. Such a dichotomy is believed to exist mainly based on the existing gap between the current upper and lower bounds for additive spanners (the current additive lower bounds match the lower bounds of its multiplicative counterpart). Perhaps surprisingly, such a dichotomy is revealed by our current results, obtained for the most basic setting of fault tolerance, namely, single edge fault and sourcewise distances.

Upper bounds for constant stretch (non-fault-tolerant) additive spanners are currently known for but a few stretch values. A $(1, 2)$ spanner with $O(n^{3/2})$ edges is presented in [1], a $(1, 6)$ spanner with $O(n^{4/3})$ edges is presented in [4], and a $(1, 4)$ spanner with $O(n^{7/5})$ edges is presented in [9]. The latter two constructions use the *path-buying* strategy, which is adopted in our additive upper bound in Sec. 4. Recently, the path-

buying strategy was employed in the context of pairwise spanners, where the objective is to construct a subgraph $H \subseteq G$ that satisfies the bounded additive stretch requirement only for a *subset* of pairs [10].

Preliminaries. Given a graph $G = (V, E)$ and a source s, let $T_0(s) \subseteq G$ be a shortest paths (or BFS) tree rooted at s. Let $\pi(x, y)$ be the (unique) $x - y$ path in $T_0(s)$. Let $E(v, G) = \{(u, v) \in E(G)\}$ be the set of edges incident to v in the graph G and let $\texttt{deg}(v, G) = |E(v, G)|$ denote the degree of vertex v in G. When the graph G is clear from the context, we may omit it and simply write $\texttt{deg}(v)$. Let $\texttt{depth}(s, v) = \text{dist}(s, v, G)$ denote the *depth* of v in the BFS tree $T_0(s)$. When the source s is clear from the context, we may omit it and simply write $\texttt{depth}(v)$ and T_0. Let $\texttt{Depth}(s) = \max_{u \in V}\{\texttt{depth}(s, u)\}$ be the *depth* of $T_0(s)$. For a subgraph $G' = (V', E') \subseteq G$ (where $V' \subseteq V$ and $E' \subseteq E$) and a pair of vertices $u, v \in V$, let $\text{dist}(u, v, G')$ denote the shortest-path distance in edges between u and v in G'. For a path $P = [u_1, \ldots, u_k]$, let $\texttt{LastE}(P)$ denote the last edge of P, let $|P|$ denote the length of P and let $P[u_i, u_j]$ be the subpath of P from u_i to u_j. For paths P_1 and P_2 where the last vertex of P_1 equals the first vertex of P_2, let $P_1 \circ P_2$ denote the path obtained by concatenating P_2 to P_1. Assuming an edge weight function $W : E(G) \to \mathbb{R}^+$, let $SP(s, u_i, G, W)$ be the set of $s - u_i$ shortest-paths in G according to the edge weights of W. (When the graph is unweighted, the parameter W is omitted.) Throughout, the edges of these paths are considered to be directed away from the source s. Given an $s - t$ path P and an edge $e = (u, v) \in P$, let $\text{dist}(s, e, P)$ be the distance (in edges) between s and e on P. In addition, for an edge $e = (u, v) \in T_0(s)$, define $\text{dist}(s, e) = i$ if $\texttt{depth}(u) = i - 1$ and $\texttt{depth}(v) = i$. For a subset $V' \subseteq V$, let $G(V')$ be the induced subgraph on V'. Let $\texttt{LCA}(V')$ be the least common ancestor of all the vertices in V'. A *replacement path* $P^*_{i,j}$ is a shortest path in $SP(s, u_i, G \setminus \{e_j\})$. Note that if $e_j \notin \pi(s, u_i)$ then the replacement path $P^*_{i,j}$ is simply the shortest-path $\pi(s, u_i)$.

Fix the source $s \in V$. For an edge $e = (u, v) \in T_0$, denote the set of vertices in $T_0(v)$, the subtree of T_0 rooted at v, by

$$\mathcal{S}(e) \;=\; V(T_0[v]).$$

Note that the vertices of $\mathcal{S}(e)$ are precisely those sensitive to the failure of the edge e, i.e., the vertices w having e on $\pi(s, w)$, their $s - w$ path in $T_0(s)$, hence also $\mathcal{S}(e) \;=\; \{w \mid e \in \pi(s, w)\}$.

2 Multiplicative FT-ABFS Structures

This section describes algorithms for constructing FT-ABFS structures for unweighted undirected graphs.

2.1 Single edge fault

We establish the following.

Theorem 2.1 *There exists a poly-time algorithm that for every n-vertex graph G and source s constructs a 1-edge $(3,0)$* `FT-ABFS` *structure with $O(n)$ edges.*

We begin by providing an informal intuition for the algorithm. The construction is based on starting from a BFS tree T_0 and adding edges to it until it satisfies the requirement. Specifically, the algorithm constructs a collection of replacement paths, $P^*_{i,j} \in SP(s, u_i, G \setminus \{e_j\})$ for every vertex-edge pair (i,j), satisfying that $e_j \in \pi(s,u_i)$. From each such path $P^*_{i,j}$, only the first new edge $e = (x,y) \in P^*_{i,j} \setminus T_0$, i.e., the new edge closest to s, is taken into the spanner.

The correctness analysis shows that the construction of the $P^*_{i,j}$ collection guarantees that the endpoint y of the first new edge e, as well as the path endpoint u_i, are both sensitive to the failure of the edge $e_j = (x', y')$, namely, $u_i, y \in T_0(y')$. Therefore, the $y - u_i$ path $\pi(y,u_i) \subseteq T_0(y') \subseteq T_0 \setminus \{e_j\}$ in the BFS tree T_0 is free of the failing edge e_j and hence provides a safe alternative path to the segment $P^*_{i,j}[y,u_i]$, which possibly might contain many edges that are missing in T_0. Then, by employing the triangle inequality, we also get that the alternative $s - u_i$ path $P^*_{i,j}[s,y] \circ \pi(y,u_i)$ is not much longer than the optimal counterpart $P^*_{i,j}$. Perhaps the more surprising part is the *size* analysis, where we show that every vertex y can appear as the endpoint of the *first* new edge of at most three replacement paths. This should be contrasted with [15], where it is shown that a vertex can be the endpoint of the *last* new edge of $\Omega(\sqrt{n})$ replacement paths. Hence, taking the last edge of every replacement path results in an exact `FT-BFS` structure with $\Theta(n^{3/2})$ edges, while taking the first new edge of every replacement path results in an approximate $(3,0)$ `FT-ABFS` structure with at most $3n$ edges.

We now provide some intuition explaining why a vertex might be the endpoint of the *first* new edge of at most *three* replacement paths. Let $P^*_{i_1,j_1}, P^*_{i_2,j_2}, \ldots, P^*_{i_k,j_k}$ be the replacement paths in which the endpoint of the first new edge is y, i.e., $\widehat{e}_\ell = \texttt{LastE}(P^*_{i_\ell,j_\ell}[s,y]) \notin T_0$ and $P^*_{i_\ell,j_\ell}[s,y] \setminus \{\widehat{e}_\ell\} \subseteq T_0$ for every $\ell \in \{1,\ldots,k\}$. We then show that upon a proper construction of the replacement paths, the fact that $\widehat{e}_\ell \notin T_0$ implies that $e_{j_1}, \ldots, e_{j_k} \subseteq \pi(s,y)$ (otherwise the original shortest path $\pi(s,y)$ could be used in $P^*_{i_\ell,j_\ell}$ instead of the segment $P^*_{i_\ell,j_\ell}[s,y]$) and letting e_{j_ℓ} be sorted in increasing distance from s on $\pi(s,y)$, it also holds that the truncated replacement paths are monotonically decreasing, i.e., $|P^*_{i_1,j_1}[s,y]| > \ldots > |P^*_{i_k,j_k}[s,y]|$. Let x_ℓ be such that $\widehat{e}_\ell = (x_\ell, y)$. Since $P^*_{i_\ell,j_\ell}[s,x_\ell] = \pi(s,x_\ell)$, we have that y is connected by an edge to k vertices $x_1,\ldots,x_k$ of distinct distances from s, $\mathrm{dist}(s,x_1,G) > \ldots > \mathrm{dist}(s,x_k,G)$, hence by the triangle inequality, necessarily $k \leq 3$.

Algorithm Description. We next formally describe the algorithm. For a path P in the constructed structure, let $\texttt{New}(P) = E(P) \setminus E(T_0)$ be the set of new edges in P, namely, edges that were added to T_0 during the construction process. Let $\texttt{FirstNewE}(P^*_{i,j})$ be the first (from s) new edge on $P^*_{i,j}$ that is not in T_0. See Fig. 1 for an illustration of these definitions.

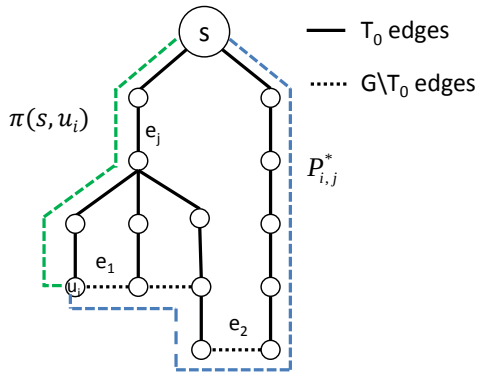

Figure 1: The path $P^*_{i,j}$ protects u_i against the failure of the edge e_j. Here $e_1 = \texttt{LastE}(P^*_{i,j})$ and $e_2 = \texttt{FirstNewE}(P^*_{i,j})$. The algorithm will add e_2 to E^*_i and subsequently to the output structure H because e_1 is also a new edge, so $P^*_{i,j}$ is new-ending.

The $(3, 0)$ `FT-ABFS` structure H is constructed by adding to T_0 only new edges that appear as the first edges on *some* of the replacement paths. The algorithm operates as follows.

Fix an ordering on the edges $E = \{e_1, \ldots, e_m\}$ and on the vertices $u_1, \ldots, u_n$. In round i, the vertex u_i is considered. The round consists of $\texttt{depth}(u_i)$ iterations. In iteration j, consider $e_j \in \pi(s, u_i)$ and define the path protecting u_i against the failure of e_j to be $P^*_{i,j} \in SP(s, u_i, G \setminus \{e_j\}, W_i)$, where $W_i : E \to \mathbb{R}_{>0}$ is a weight assignment for the edges of E defined by

$$W_i(e_\ell) = \begin{cases} \omega_\ell, & \text{if } e_\ell \in E(T_0) \setminus E(\pi(s, u_i)), \\ \omega_\ell + \epsilon_2, & \text{if } e_\ell \in \pi(s, u_i), \\ \omega_\ell + \epsilon_1, & \text{otherwise,} \end{cases} \tag{1}$$

where $\omega_\ell = n^6 \cdot 2^{m+1} + 2^\ell$, $\epsilon_1 = n^3 \cdot 2^{m+1}$ and $\epsilon_2 = 2^{m+1}$. Call a replacement path $P^*_{i,j}$ *new-ending* if its last edge is new, namely, $\texttt{LastE}(P^*_{i,j}) \notin E(T_0)$. For every vertex $u_i \in V$, define

$$E^*_i \;=\; \{\texttt{FirstNewE}(P^*_{i,j}) \mid \; e_j \in \pi(s, u_i) \text{ and } P^*_{i,j} \text{ is new-ending}\}.$$

Let $E^*_{new} = \bigcup_{u_i \in V} E^*_i$ and $H = T_0 \cup E^*_{new}$.

Correctness. We now prove the correctness of the algorithm and establish Thm. 2.1, by showing that taking into the constructed H merely the first new edge from each new-ending replacement path $P^*_{i,j}$ is sufficient in order to guarantee the existence of an approximate $s-u_i$ replacement path in the surviving structure $H \setminus \{e_j\}$, for every u_i and $e_j \in \pi(s,u_i)$.

Let us start by explaining the specific weight assignment chosen. The role of $W_i(e_\ell)$ is to enforce a unique $s-u_i$ shortest-path in $SP(s,u_i,G\setminus\{e_j\})$. This is important for both the correctness and the size analysis of the $(3,0)$ `FT-ABFS` structure. (In other words, carelessly taking the first edge of an arbitrary replacement path might result in a dense subgraph which is also not a $(3,0)$ `FT-ABFS` structure.) The weight assignment achieves this as follows. For every $u_i \in V$, let $\texttt{Cost}_i(P) = \sum_{e\in P} W_i(e)$ be the weighted cost of P, i.e., the sum of its edge weights. Then given paths P_1, P_2, the weight assignment W_i has the following properties, implying that $\texttt{Cost}_i(P)$ can in some sense be viewed as based on $|P|$, $|\texttt{New}(P)|$ and $|E(P)\cap\pi(s,u_i)|$ lexicographically.

Fact 2.2 *For every two paths P_1, P_2 and $u_i \in V$,*

(a) If $|P_1| < |P_2|$, then $\texttt{Cost}_i(P_1) < \texttt{Cost}_i(P_2)$.

(b) If $|P_1| = |P_2|$ and $|\texttt{New}(P_1)| < |\texttt{New}(P_2)|$, then $\texttt{Cost}_i(P_1) < \texttt{Cost}_i(P_2)$.

(c) If $|P_1| = |P_2|$, $|\texttt{New}(P_1)| = |\texttt{New}(P_2)|$ and $|E(P_1)\cap E(\pi(s,u_i))| < |E(P_2)\cap E(\pi(s,u_i))|$, then $\texttt{Cost}_i(P_1) < \texttt{Cost}_i(P_2)$.

(d) If $|P_1| = |P_2|$, $|\texttt{New}(P_1)| = |\texttt{New}(P_2)|$ and $|E(P_1)\cap E(\pi(s,u_i))| = |E(P_2)\cap E(\pi(s,u_i))|$, then $\texttt{Cost}_i(P_1) < \texttt{Cost}_i(P_2)$ iff $\sum_{e_k\in P_1}\omega_k < \sum_{e_k\in P_2}\omega_k$.

Conversely we also have the following.

Fact 2.3 *If $\texttt{Cost}_i(P_1) < \texttt{Cost}_i(P_2)$, then necessarily one of the following four conditions holds:*

(a) $|P_1| < |P_2|$,

(b) $|\texttt{New}(P_1)| < |\texttt{New}(P_2)|$,

(c) $|E(P_1)\cap E(\pi(s,u_i))| < |E(P_2)\cap E(\pi(s,u_i))|$,

(d) $\sum_{e_k\in P_1}\omega_k < \sum_{e_k\in P_2}\omega_k$.

The following key observation is used repeatedly in what follows.

Observation 2.4 *For every replacement path $P^*_{i,j}$ and every new edge $e=(x,y) \in \texttt{New}(P^*_{i,j})$ on it, $y \in \mathcal{S}(e_j)$ (or $e_j \in \pi(s,y)$).*

Proof: Assume, towards contradiction, that $e = (x, y) \in \texttt{New}(P^*_{i,j})$ and yet $e_j \notin \pi(s, y)$. Since $e = (x, y) \in \texttt{New}(P^*_{i,j})$ is a new edge, it holds that $P^*_{i,j}[s, y] \neq \pi(s, y)$. Consider an alternative $s - u_i$ replacement path $P' = \pi(s, y) \circ P^*_{i,j}[y, u_i]$. Since $|\pi(s, y)| \leq |P^*_{i,j}[s, y]|$ but $|\texttt{New}(\pi(s, y))| = 1$ and $|\texttt{New}(P^*_{i,j}[s, y])| \geq 1$ (since T_0 contains only one $s - y$ path, $\pi(s, y)$). It follows by Obs. 2.2, that $\texttt{Cost}_i(\pi(s, y)) < \texttt{Cost}_i(P^*_{i,j}[s, y])$ and thus also $\texttt{Cost}_i(P') < \texttt{Cost}_i(P^*_{i,j})$, in contradiction to the fact that $P^*_{i,j} \in SP(s, u_i, G \setminus \{e_j\}, W_i)$. The observation follows. ∎

We next provide the following claim, showing that if a BFS edge $e_j \in \pi(s, u_i)$ fails, then adding the first new edge e' of a new-ending replacement path $P^*_{i,j}$ to the BFS tree T_0, recovers its connectivity. In order words, $T' = T_0 \setminus \{e_j\} \cup \{e'\}$ is a connected spanning tree of the graph G. Moreover, we show that the $s - u_i$ path in T' has low stretch compared to the $s - u_i$ shortest-path in $G \setminus \{e_j\}$.

Lemma 2.5 *Let $P^*_{i,j}$ be a new-ending replacement path, let $e_0 = (w_1, w_2) = \texttt{FirstNewE}(P^*_{i,j})$ and let $T' = T_0 \cup \{e_0\}$. Then* $\text{dist}(s, u_i, T' \setminus \{e_j\}) \leq 3 \cdot |P^*_{i,j}|$.

Proof: Let $e_j = (y_1, y_2)$; see Fig. 2 for illustration. Note that by Obs. 2.4, w_2 is in $\mathcal{S}(e_j)$, hence both $u_i, w_2 \in T_0(y_2)$, the subtree of T_0 rooted at y_2. Therefore the path between w_2 and u_i, $\pi(w_2, u_i)$, does not use e_j, hence it exists in $T_0 \setminus \{e_j\}$.

Let $x = \texttt{LCA}(u_i, w_2)$ be the least common ancestor of u_i and w_2 in T_0. Then $x \in \pi(y_2, u_i)$. Let $A = P^*_{i,j}[s, w_2], B = \pi(s, x), C = \pi(x, w_2), D = P^*_{i,j}[w_2, u_i]$ and $R = \pi(x, u_i)$. Consider an alternative $s - u_i$ replacement path $P' = A \circ C \circ R$ that uses the $w_2 - u_i$ path in T_0 (see Fig. 2). Note that since $e_0 = (w_1, w_2)$ is the first new edge on $P^*_{i,j}$, it follows that $P' \subseteq T' = T_0 \cup \{e_0\}$. Since P' is a replacement path for u_i in T', it remains to bound its length. Note that $|P'| = |A| + |C| + |R|$ and $|P^*_{i,j}| = |A| + |D|$. First, since A is a shortest $s - w_2$ path in $G \setminus \{e_j\}$ but $B \circ C$ is an $s - w_2$ shortest-path in G, it follows that $|B| + |C| \leq |A|$. Next, consider the two $x - u_i$ paths R and $C \circ D$. Since R is a shortest $x - u_i$ path, it follows that $|R| \leq |C| + |D|$. Therefore $|P'| \leq |A| + 2|C| + |D| \leq 3|A| + |D| \leq 3|P^*_{i,j}|$. The lemma follows. ∎

Lemma 2.6 *H is a $(3, 0)$ $\texttt{FT-ABFS}$ structure.*

Proof: Assume, towards contradiction, that H is not a $(3, 0)$ $\texttt{FT-ABFS}$ structure. Let $BP = \{(i, j) \mid \text{dist}(s, u_i, H \setminus \{e_j\}) > 3|P^*_{i,j}|\}$ be the set of "bad pairs," namely, vertex-edge pairs (i, j) for which the length of the replacement $s - u_i$ path in $H \setminus \{e_j\}$ is greater than $3 \cdot \text{dist}(s, u_i, G \setminus \{e_j\})$. (By the contradictory assumption, $BP \neq \emptyset$.) For each bad pair $(i, j) \in BP$, define $BE(i, j) = P^*_{i,j} \setminus E(H)$ to be the set of "bad edges," namely, the set of $P^*_{i,j}$ edges that are missing in H. By definition, $BE(i, j) \neq \emptyset$ for every bad pair $(i, j) \in BP$. Let $d(i, j) = \max_{e \in BE(i,j)}\{\text{dist}(s, e, P^*_{i,j})\}$ be the maximal depth of a missing edge in $BE(i, j)$, and let $DM(i, j)$ denote that "deepest missing edge", i.e., the

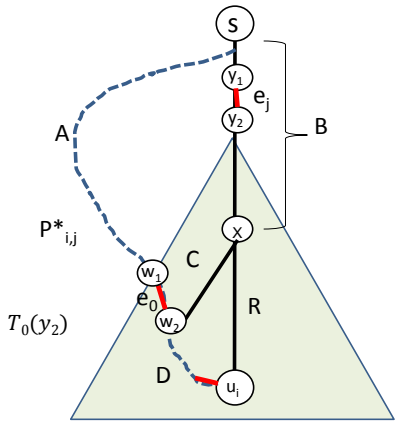

Figure 2: Illustration of the approximate replacement path. Solid lines represent tree edges.

edge e on $P^*_{i,j}$ satisfying $d(i,j) = \text{dist}(s, e, P^*_{i,j})$. Finally, let $(i',j') \in BP$ be the pair that minimizes $d(i,j)$, and let $e_1 = (v_{\ell_1}, u_{i_1}) \in BE(i',j')$ be the deepest missing edge on $P^*_{i',j'}$, namely, $e_1 = DM(i',j')$. Note that e_1 is the *shallowest* "deepest missing edge" over all bad pairs $(i,j) \in BP$. By Obs. 2.4, $e_{j'} \in \pi(s, u_{i_1})$.

Consider the $s - u_{i_1}$ replacement path $P^*_{i_1,j'}$. Note that there are two $s - u_{i_1}$ replacement paths, $P_1 = P^*_{i',j'}[s, u_{i_1}]$ and $P_2 = P^*_{i_1,j'} \in G \setminus \{e_{j'}\}$, and by their optimality we have that $|P_1| = |P_2|$ (these paths might - but do not have to - be the same).

We distinguish between two cases: (C1) $\texttt{LastE}(P_2) \notin T_0$ and (C2) $\texttt{LastE}(P_2) \in T_0$. Begin with case (C1). By construction, $\texttt{FirstNewE}(P_2) \in E^*_{i_1}$, so $\texttt{FirstNewE}(P_2) \in H$. By Lemma 2.5, there exists an $s - u_{i_1}$ replacement path P' in $G \setminus \{e_{j'}\}$ such that $|P'| \leq 3 \cdot |P_2|$ and $P' \subseteq (T_0 \cup \{\texttt{FirstNewE}(P_2)\}) \subseteq H$. Consider the $s - u_{i'}$ replacement path

$$P'' \;=\; P' \circ P^*_{i',j'}[u_{i_1}, u_{i'}].$$

Note that since e_1 is the deepest missing edge in $P^*_{i',j'}$, it holds that $P^*_{i',j'}[u_{i_1}, u_{i'}] \subseteq H$ and by the previous argument $P' \subseteq H \setminus \{e_{j'}\}$, concluding that P'' is an $s - u_{i'}$ replacement path in $H \setminus \{e_{j'}\}$ Moreover, its length is bounded by

$$|P''| \;=\; |P'| + |P^*_{i',j'}[u_{i_1}, u_{i'}]| \leq 3|P_1| + |P^*_{i',j'}[u_{i_1}, u_{i'}]| \;\leq\; 3|P^*_{i_1,j'}|,$$

contradicting the fact that $(i',j') \in BP$ is a bad pair.

Now consider case (C2) where $\texttt{LastE}(P_2) \in T_0$. We show that in this case $(i_1,j') \notin BP$. Assume, towards contradiction, that (i_1,j') is a bad pair. This implies that $P_2 \nsubseteq H$. Since $|P_1| = |P_2|$, $\texttt{LastE}(P_1) \notin T_0$ but $\texttt{LastE}(P_2) \in T_0$, it holds that the "deepest missing edge" e'' in P_2 is such that $\text{dist}(s, e'', P_2) < \text{dist}(s, e_1, P_1)$ (or $d(i_1,j') < d(i',j')$) in contradiction to the selection of (i',j'). Hence, we conclude that $(i_1,j') \notin BP$,

which guarantees the existence of an $s - u_{i_1}$ replacement path $P' \in G \setminus \{e_{j'}\}$ such that $|P'| \leq |P_2|$. Finally, the $s - u_{i'}$ path $P'' = P' \circ P^*_{i',j'}[u_{i_1}, u_{i'}]$ exists in $H \setminus \{e_{j'}\}$ and $|P''| \leq 3|P_1| + |P^*_{i',j'}[u_{i_1}, u_{i'}]| \leq 3|P^*_{i',j'}|$, in contradiction to the fact that $(i', j') \in BP$. The lemma follows. ∎

Size analysis.

Lemma 2.7 $|E(H) \setminus E(T_0)| = |E^*_{new}| \leq 3n$.

Proof: We show that every vertex u_i can have at most 3 of its incident edges in E^*_{new}. Assume, towards contradiction, that there exists some u_i with (at least) 4 edges in E^*_{new}, $e_k = (v_k, u_i)$ for $k \in \{1, \ldots, 4\}$, that appear as first new edges in the replacement paths $P^*_{i_1,j_1}, P^*_{i_2,j_2}, P^*_{i_3,j_3}, P^*_{i_4,j_4}$ respectively. By Obs. 2.4, it holds that the 4 failed edges $e_{j_k} \in \pi(s, u_i)$, $k \in \{1, \ldots, 4\}$ appear on $\pi(s, u_i)$ and by definition, $e_{j_k} \in \pi(s, u_{i_k})$, for every $k \in \{1, \ldots, 4\}$. Without loss of generality, assume that $\text{dist}(s, e_{j_k}, \pi(s, u_i)) \leq \text{dist}(s, e_{j_{k+1}}, \pi(s, u_i))$ for every $k \in \{1, 2, 3\}$, namely, that the edges $e_{j_1}, e_{j_2}, e_{j_3}, e_{j_4}$ occur on $\pi(s, u_i)$ in that order. For illustration see Fig. 3.

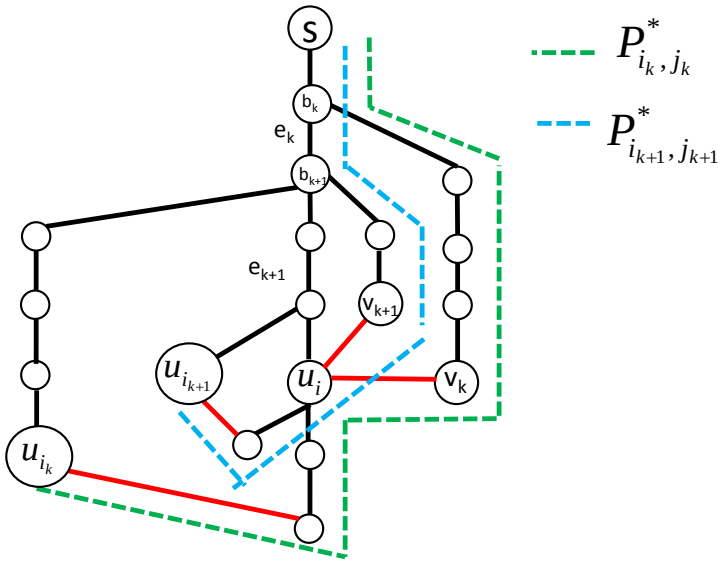

Figure 3: Illustration of the replacement paths in which edge of u_i are the first new edges. New edges are represented in red.

Consider the 4 truncated $s - u_i$ paths $P_k = P^*_{i_k,j_k}[s, u_i]$ for $k = \{1, \ldots, 4\}$. Note that since $\texttt{LastE}(P_k) = e_k$ is the only new edge in P_k, i.e., $P_k \setminus \{\texttt{LastE}(P_k)\}$ has no new edges, or,

$$P_k \setminus \{\texttt{LastE}(P_k)\} \subseteq T_0 \setminus \{e_{j_k}\}. \tag{2}$$

Let b_k be the first divergence point of P_k from $\pi(s,u_i)$, namely, the last vertex on that $\pi(s,u_i)$ for which $\pi(s,b_k) = P_k[s,b_k]$. Let $Q_{Pref}[k] = \pi[s,b_k] = P_k[s,b_k]$ (where the equality is by the definition of b_k) be the maximal common prefix of the paths P_k and $\pi(s,u_{i_k})$, for $k \in \{1,\ldots,4\}$. When k is clear from the context, we may omit it and simply write Q_{Pref}. Let $Q_{Suff}[k] = P_k[b_k,u_i]$. We now show that b_k is the *only* divergence point of P_k and $\pi(s,u_i)$, or in other words, the paths meet again only at u_i. Formally, we show the following.

Claim 2.8 *$(V(Q_{Suff}[k]) \cap V(\pi(s,u_i))) \setminus \{b_k,u_i\} = \emptyset$ for $k \in \{1,\ldots,4\}$.*

Proof: Assume, towards contradiction, that the paths intersect again at some vertex

$$w \in (V(Q_{Suff}[k]) \cap V(\pi(s,u_i))) \setminus \{b_k,u_i\}.$$

Recall that u_{i_k} was considered in round i_k and let j_k be the iteration in this round in which the edge j_k was considered. Let $Q_0 = Q_{Pref}[k], Q_1 = P_k[b_k,w], Q'_1 = \pi(b_k,w), Q_2 = P_k[w,u_i], Q'_2 = \pi(w,u_i)$. For illustration, see Fig. 4. We distinguish between two cases concerning the faulty edge e_{j_k}: (C1) $e_{j_k} \in Q'_1$ or (C2) $e_{j_k} \in Q'_2$.

In case (C1), $Q'_2 \subseteq G \setminus \{e_{j_k}\}$, and as it is part of the BFS tree, it holds that $|Q'_2| \leq |Q_2|$. Since Q'_2 is free of new edges but $\texttt{LastE}(Q_2) = e_k$ is new, it holds that $\texttt{Cost}_{i_k}(Q'_2) < \texttt{Cost}_{i_k}(Q_2)$, in contradiction to the fact that $P^*_{i_k,j_k} \in SP(s,u_{i_k},G \setminus \{e_{j_k}\},W_{i_k})$.

In case (C2), $Q'_1 \subseteq G \setminus \{e_{j_k}\}$, and as it is part of the BFS tree, it holds that $|Q'_1| \leq |Q_1|$. Since T_0 contains a single $b_k - w$ path corresponding to Q'_1, it must hold that $Q_1 \not\subseteq T_0$ (Q_1 has at least one new edge) and therefore $\texttt{Cost}_{i_k}(Q'_1) < \texttt{Cost}_{i_k}(Q_1)$, in contradiction again to the fact that $P^*_{i_k,j_k} \in SP(s,u_{i_k},G \setminus \{e_{j_k}\},W_{i_k})$. ∎

It follows from Cl. 2.8, that $E(P_k) \cap E(\pi(s,u_i)) = Q_{Pref}[k]$. We now focus on the edge-set intersections

$$\mathcal{I}_{k,\ell} \;=\; E(P_k) \cap E(\pi(s,u_{i_\ell}))$$

and establish the following auxiliary claim, showing that the same holds also for the complete path $\pi(s,u_{i_\ell})$, for the ℓ values needed later.

Claim 2.9 *$\mathcal{I}_{k,\ell} = Q_{Pref}[k]$*
(a) for every $\ell \in \{k,\ldots,4\}$, and
(b) for every $\ell < k$ such that $\texttt{depth}(b_k) \leq \texttt{depth}(b_\ell)$.

Proof: Recall that $P_k = Q_{Pref}[k] \circ Q_{Suff}[k]$ and let $e_{j_k} = (x_k,y_k)$ and $e_{j_\ell} = (x_\ell,y_\ell)$. We prove parts (a) and (b) in two steps. We first show that $Q_{Pref}[k] \subseteq \pi(s,u_{i_\ell})$ and then show that $Q_{Suff}[k]$ and $\pi(s,u_{i_\ell})$ are edge disjoint for ℓ satisfying (a) or (b). We begin by showing that $Q_{Pref}[k] = \pi(s,b_k) \subseteq \pi(s,u_{i_\ell})$. Let $\ell \in \{k,\ldots,4\}$. Since $e_{j_\ell} \in \pi(s,u_{i_\ell})$, by the ordering of the edges e_{j_k}, it holds that also $e_{j_k} = (x_k,y_k) \in \pi(s,u_{i_\ell})$. Since

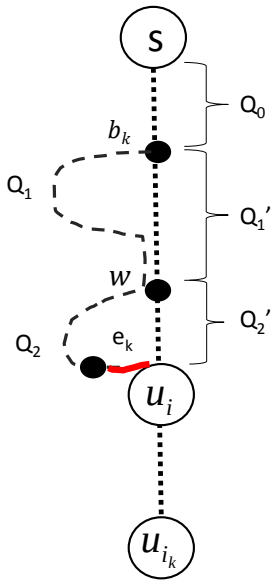

Figure 4: The dotted straight line is $\pi(s, u_i)$ and the dashed line depicts $Q_{Suff}[k] = P_k[b_k, u_i]$. The divergence point is unique. The red e_k is the new edge of u_i. The faulty edge e_{j_k} can occur in either Q_1' or Q_2'.

$P_k \subseteq G \setminus \{e_{j_k}\}$, the divergence point b_k of P_k and $\pi(s, u_i)$ occurred above y_k, hence $\pi(s, b_k) \subseteq \pi(s, u_{i_\ell})$. Next, let $\ell < k$ be such that $\texttt{depth}(b_k) \leq \texttt{depth}(b_\ell)$. By part (a), $\pi(s, b_\ell) \subseteq \pi(s, u_{i_\ell})$. Since the divergence point b_k occurred not after b_ℓ (and both b_k and b_ℓ are in $\pi(s, u_i)$), it holds that also $\pi(s, b_k) \subseteq \pi(s, u_{i_\ell})$.

Next, we consider $Q_{Suff}[k]$ and show that it is edge disjoint from $\pi(s, u_{i_\ell})$ for $\ell \in \{k, \ldots, 4\}$. By the above argumentation, $\pi(s, u_{i_\ell}) = \pi(s, y_k) \circ \pi(y_k, u_{i_\ell})$. By Cl. 2.8, the paths $Q_{Suff}[k]$ and $\pi(s, u_i)$ are edge disjoint and hence the two paths $\pi(s, y_k)$ and $Q_{Suff}[k]$ are edge disjoint. It remains to show that $\pi(y_k, u_{i_\ell})$ and $Q_{Suff}[k]$ are edge disjoint. Since by Eq. (2) the path $P' = P_k \setminus \{\texttt{LastE}(P_k)\}$ exists in $T_0 \setminus \{e_{j_k}\}$, it holds that $V(P') \cap \mathcal{S}(e_{j_k}) = \emptyset$. However, $V(\pi(y_k, u_{i_\ell})) \subseteq \mathcal{S}(e_{j_k})$. Hence P' does not intersect with $\pi(s, u_{i_\ell})$. Finally, let $\ell < k$ be as in (b), i.e., such that $\texttt{depth}(b_k) \leq \texttt{depth}(b_\ell)$. By Cl. 2.8, $Q_{Suff}[k]$ and $\pi(s, u_i)$ are edge disjoint and hence $\pi(s, y_\ell) \subseteq \pi(s, u_i)$ and $Q_{Suff}[k]$ are edge disjoint. It remains to show that $\pi(y_\ell, u_{i_\ell})$ and $Q_{Suff}[k]$ are edge disjoint. Since b_k diverged from $\pi(s, u_i)$ not after b_ℓ, it holds that b_k is above y_ℓ hence $P' \subseteq T_0 \setminus \{e_{j_\ell}\}$. Therefore $V(P') \cap \mathcal{S}(e_{j_\ell}) = \emptyset$, but $V(\pi(y_\ell, u_{i_\ell})) \subseteq \mathcal{S}(e_{j_\ell})$, hence $\pi(y_\ell, u_{i_\ell})$ and $Q_{Suff}[k]$ are edge disjoint as required. The claim follows. ∎

Claim 2.10 $\texttt{depth}(b_k, T_0) < \texttt{depth}(b_{k+1}, T_0)$ *for every* $k \in \{1, 2, 3\}$.

Proof: Towards contradiction, assume that $\texttt{depth}(b_k, T_0) \geq \texttt{depth}(b_{k+1}, T_0)$ for some

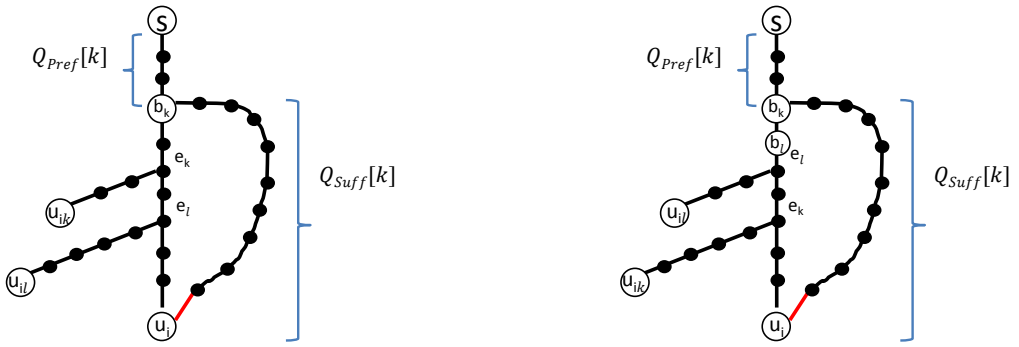

(a) (b)

Figure 5: Illustration for Cl. 2.9. Black edges represent the original BFS tree T_0. The red edge is a new edge.

$k \in \{1,2,3\}$. We first claim that in this case both $P_{k+1}, P_k \subseteq G \setminus \{e_{j_k}, e_{j_{k+1}}\}$. Recall that $P_k = Q_{Pref}[k] \circ Q_{Suff}[k]$ and $P_{k+1} = Q_{Pref}[k+1] \circ Q_{Suff}[k+1]$. Note that by Cl. 2.8, the paths $Q_{Suff}[k], Q_{Suff}[k+1]$ are edge disjoint with $\pi(s,u_i)$ and therefore $Q_{Suff}[k], Q_{Suff}[k+1] \subseteq G \setminus \{e_{j_k}, e_{j_{k+1}}\}$. In addition, by the fact that P_k diverged from $\pi(s,u_i)$ before the faulty edge e_{j_k} (and by ordering also before $e_{j_{k+1}}$) it holds that $Q_{Pref}[k] \subseteq G \setminus \{e_{j_k}, e_{j_{k+1}}\}$. Since $\texttt{depth}(b_k, T_0) \geq \texttt{depth}(b_{k+1}, T_0)$, it holds that P_{k+1} diverged from $\pi(s,u_i)$ not after b_k, hence $Q_{Suff}[k], Q_{Suff}[k+1] \subseteq G \setminus \{e_{j_k}, e_{j_{k+1}}\}$. Overall, we get that $P_{k+1}, P_k \subseteq G \setminus \{e_{j_k}, e_{j_{k+1}}\}$.

We thus have two alternative replacement $s - u_{i_k}$ (resp., $s - u_{i_{k+1}}$) paths given by $\widetilde{P}_k = P_{k+1} \circ P^*_{i_k,j_k}[u_i, u_{i_k}]$ and $\widetilde{P}_{k+1} = P_k \circ P^*_{i_{k+1},j_{k+1}}[u_i, u_{i_{k+1}}]$ respectively. We now derive a contradiction by analyzing the costs of $P^*_{i_k,j_k}$ and $P^*_{i_{k+1},j_{k+1}}$ and showing that all costs components (see Eq. (1)) are equal except the last. By the optimality of $P^*_{i_k,j_k}$ and $P^*_{i_{k+1},j_{k+1}}$ in round i_k and i_{k+1} respectively, i.e., by the fact that $P^*_{i_k,j_k} \in SP(s, u_{i_k}, G \setminus \{e_{j_k}\}, W_{i_k})$ and $P^*_{i_{k+1},j_{k+1}} \in SP(s, u_{i_{k+1}}, G \setminus \{e_{j_{k+1}}\}, W_{i_{k+1}})$, it follows that $|P_k| = |P_{k+1}|$. In addition, since P_k is a subpath of $P^*_{i_k,j_k}$ and its last edge is the first new edge of $P^*_{i_k,j_k}$ (i.e., $\texttt{LastE}(P_k) = \texttt{FirstNewE}(P^*_{i_k,j_k})$), it follows that $|\texttt{New}(P_k)| = 1$. By a similar argument, since $\texttt{LastE}(P_{k+1}) = \texttt{FirstNewE}(P^*_{i_{k+1},j_{k+1}})$, we also have that $|\texttt{New}(P_{k+1})| = 1$. By the optimality of $P^*_{i_k,j_k} \in SP(s, u_{i_k}, G \setminus \{e_{j_k}\}, W_{i_k})$ according to weight assignments W_{i_k} (see Fact 2.2(c)) we get that

$$|\mathcal{I}_{k,k}| \leq |\mathcal{I}_{k+1,k}| \ . \tag{3}$$

In the same manner, by the optimality of $P^*_{i_{k+1},j_{k+1}} \in SP(s, u_{i_{k+1}}, G \setminus \{e_{j_{k+1}}, W_{i_{k+1}}\})$ according to weight assignments $W_{i_{k+1}}$, we get that

$$|\mathcal{I}_{k+1,k+1}| \leq |\mathcal{I}_{k,k+1}| \ . \tag{4}$$

Applying Cl. 2.9(a) with $\ell = k$ we have that

$$|\mathcal{I}_{k,k}| = |Q_{Pref}[k]| = \texttt{depth}(b_k) \ . \tag{5}$$

Applying Cl. 2.9(a) with $\ell = k+1$ we have that

$$|\mathcal{I}_{k,k+1}| = |Q_{Pref}[k]| = \texttt{depth}(b_k) \tag{6}$$

and

$$|\mathcal{I}_{k+1,k+1}| = |Q_{Pref}[k+1]| = \texttt{depth}(b_{k+1}) \ . \tag{7}$$

By. Cl. 2.9(b), we also have that

$$|\mathcal{I}_{k+1,k}| = |Q_{Pref}[k+1]| = \texttt{depth}(b_{k+1}) \ . \tag{8}$$

Combining Eq. (3) with Eq. (5) and (8), we get that $\texttt{depth}(b_k) \leq \texttt{depth}(b_{k+1})$. Combining Eq. (4) with Eq. (6) and (7), we get the opposite inequality, $\texttt{depth}(b_{k+1}) \leq \texttt{depth}(b_k)$. It follows that $\texttt{depth}(b_k) = \texttt{depth}(b_{k+1})$, hence inequalities (3) and (4) are in fact equalities.

As we have shown that the paths $P^*_{i_k,j_k}$ and $P^*_{i_{k+1},j_{k+1}}$ have the same length, the same number of new edges and the same number of joint edges with the shortest-path, by Fact 2.2(d) their relative costs are determined by $\widetilde{W}_k = \sum_{e_j \in P_k} \omega_j$ and $\widetilde{W}_{k+1} = \sum_{e_j \in P_{k+1}} \omega_j$. Hence, by the optimality of $P^*_{i_k,j_k}$ under $\texttt{Cost}_{i_k}$ it follows that $\widetilde{W}_k < \widetilde{W}_{k+1}$, and by the optimality of $P^*_{i_{k+1},j_{k+1}}$ under $\texttt{Cost}_{i_{k+1}}$ we get that $\widetilde{W}_{k+1} < \widetilde{W}_k$, contradiction. ∎

Claim 2.10 implies that the vertices b_1, b_2, b_3, b_4 are distinct, and moreover, they appear in this order on $\pi(s, u_i)$. In addition, note that for every $k \in \{1,2,3\}$, $b_k \in P_{k+1}$ (since b_{k+1} is below b_k on $\pi(s,u_i)$).

Claim 2.11 *$|P_k[b_k, u_i]| > |P_{k+1}[b_k, u_i]|$ for every $k \in \{1,2,3\}$.*

Proof: By the uniqueness of the divergence point b_k of $\pi(s,u_i)$ and P_k (Cl. 2.8) and by Cl. 2.10, $P_k \subseteq G \setminus \{e_{j_k}, \ldots, e_{j_4}\}$ for every $k \in \{1,2,3,4\}$. Since P_k is a replacement path in $G \setminus \{e_{j_{k'}}\}$ for $k' > k$, but P_k was nevertheless not chosen as part of the replacement path $P^*_{i_{k+1},j_{k+1}}$, it follows that $\texttt{Cost}_{i_{k+1}}(P_{k+1}) < \texttt{Cost}_{i_{k+1}}(P_k)$. Let us now analyze which cost component accounts for this difference. By Cl. 2.9(a), $|\mathcal{I}_{k+1,k+1}| = \texttt{depth}(b_{k+1})$ and $|\mathcal{I}_{k,k+1}| = \texttt{depth}(b_k)$. Hence, since due to Cl. 2.10, $\texttt{depth}(b_{k+1}) > \texttt{depth}(b_k)$ (and thus $|\mathcal{I}_{k+1,k+1}| > |\mathcal{I}_{k,k+1}|$) and $|\texttt{New}(P_k)| = |\texttt{New}(P_{k+1})| = 1$, it follows by the optimality of $P^*_{i_{k+1},j_{k+1}} \in SP(s, u_{i_{k+1}}, G \setminus \{e_{j_{k+1}}\}, W_{i_{k+1}})$ for i_{k+1} (see Fact 2.2(c)) that $|P_{k+1}| < |P_k|$. As $P_k[s,b_k] = P_{k+1}[s,b_k] = \pi(s,b_k) = Q_{Pref}[k]$, the claim follows. ∎

Let $d_0 = \texttt{depth}(u_i)$ and let $d_k = \texttt{depth}(v_k)$ for $k = \{1, \ldots, 4\}$. Combining Cl. 2.10 and 2.11, we get that $d_1 > d_2 > d_3 > d_4$, and as these are integers, necessarily $d_1 - d_4 \geq 3$. Now consider the edges $e_k = (v_k, u_i) = \texttt{LastE}(P_k)$ for $k = \{1, \ldots, 4\}$. The existence of these edges in G, implies that $d_k \in \{d_0 - 1, d_0, d_0 + 1\}$ for every $k = \{1, \ldots, 4\}$, implying that $d_1 - d_4 \leq 2$, contradiction. Lemma 2.7 follows. ∎

2.2 Multiple edge faults

In this section, we consider the case of f edge failures for constant $f \geq 1$, and establish the following.

Theorem 2.12 *There exists a poly-time algorithm that for every n-vertex graph constructs*

(1) a $(3(f+1),(f+1)\log n)$ FT-ABFS structure with $O(fn)$ edges and

(2) a $(3(f+1)+1,0)$ FT-ABFS structure with $O(fn+n^{1+1/k}+n\cdot((f+1)\cdot(2k-1))^{f+1})$ edges,

overcoming up to f edge faults, for every $k \geq 3$.

For an edge set $F \subseteq E$, let $P^*_F(u_i) \in SP(s,u_i,G\setminus F)$ be the $s-u_i$ replacement path upon the failure of F in G. To avoid complications due to shortest-paths of the same length, we assume all shortest-path are computed with a weight assignment W that guarantees the uniqueness of the shortest-paths.

Algorithm Description. The algorithm consists of three phases. The first phase constructs a (possibly dense) f-edge FT-BFS structure T_1 with respect to s by using Alg. ConsFT-BFS(s,G,f) to be defined later. The second phase constructs an f-edge $(3(f+1),0)$ FT-ABFS structure $T_2 \subseteq T_1$ by carefully sparsifying the edges of T_1. However, T_2 might still be dense. Finally, the last phase obtains a *sparse* $(3(f+1),(f+1)\cdot\log n)$ FT-ABFS structure $H \subseteq T_2$ where $|H| = O(n)$. We rely on the following fact.

Lemma 2.13 *[8] There exists an algorithm* ConsSpan(G,α,f) *that given an n-vertex graph G constructs an f* edge *fault tolerant $(\alpha,0)$ spanner $G' \subseteq G$ such that $|G'| \leq O(f\cdot n^{1+1/\alpha})$.*

Algorithm ConsFT-AddSpan(s,G,f) **- overview**

(1) Invoke Alg. ConsFT-BFS(s,G,f) to generate an f-edge FT-BFS structure T_1 with respect to s.

(2) Sparsify the new edges of T_1 to obtain an $(3(f+1),0)$ FT-ABFS structure $T_2 \subseteq T_1$.

(3) Set $T' \leftarrow$ ConsSpan$(T_2 \setminus T_0, \log n, f)$.

(4) $H = T_0 \cup T'$.

Algorithm $\texttt{ConsFT-BFS}(s, G, f)$ of phase (1) operates in a "brute-force" manner. For every $u_i \in V$, and every subset $F \subseteq E$, $|F| \leq f$, it constructs an $s - u_i$ replacement path $P^*_F(u_i) \in SP(s, u_i, G \setminus F, W)$ of minimal length. The f-edge $\texttt{FT-BFS}$ structure T_1 is then given by $T_1 = \{P^*_{i,F} \in SP(s, u_i, G \setminus F) \mid F \subseteq E, |F| \leq f, u_i \in V\}$. In phase (2), each of these replacement paths $P = P^*_F(u_i)$ is considered and at most $f + 1$ of the set of new edges $\texttt{New}(P)$ are taken into T_2 at the expense of introducing a stretch. To choose these edges from $\texttt{New}(P)$ for a given replacement path $P = P^*_F(u_i)$, the algorithm labels the vertices $V(P)$ according to their sensitivity to the set of failed edges F, where vertices are given the same label iff they appear in the same connected tree in the surviving forest $T_0 \setminus F$. Since vertices u', u'' of the same label remain connected in $T_0 \setminus F$, the algorithm exploits their path in $T_0 \setminus F$ as a replacement to the path $P[u', u'']$ used by the optimal replacement path. The benefit of these bypasses is that they use only edges of the original BFS, allowing us to save the new edges that occur on $P[u', u'']$. The potential drawback is that these bypasses might be longer than their counterparts in P. In the analysis we argue that the stretch introduced by these replacements is bounded by $3(f + 1)$. Hence phase (2) turns the *exact* FT-BFS structure T_1 to an *approximate* FT-BFS structure T_2 with bounded multiplicative stretch. We now describe formally the construction of T_2, beginning with the labeling scheme $\texttt{Label}_F : V \to \mathcal{C}(T_0 \setminus F)$. Let $\mathcal{C}(T_0 \setminus F) = \{T_0^1, \ldots, T_0^\ell\}$ be the set of connected components (subtrees) of the forest $T_0 \setminus F$. For pictorial illustration, see Fig. 6. Then the label $\texttt{Label}_F(u')$ of every vertex $u' \in P$ is set to $\ell' \in \{1, \ldots, \ell\}$ iff $u' \in T_0^{\ell'}$. The procedure for selecting at most $f + 1$ new edges of P using the labels $\texttt{Label}_F(u')$, $u' \in P$ is as follows. Let $\texttt{New}(P) = \{(u_1, v_1), \ldots, (u_k, v_k)\}$ be the sorted set of new edges $E(P) \setminus E(T_0)$ according to their order of appearance on the replacement path P (from s). Pair the vertices v_i by matching v_i with the farthest $v_{i'}$ for $i' \in \{i, \ldots, k\}$ of the same label $\texttt{Label}_F(v_i) = \texttt{Label}_F(v_{i'})$, setting $M(v_i) = v_{i'}$ and $\widehat{m}(v_i) = i'$. Initialize $i = 1$ and $T_2 = T_0$. Now, repeatedly (until $i \geq k$) add the new edge $e_i = (u_i, v_i)$ to T_2 and set $i = \widehat{m}(v_i) + 1$.

Finally, we describe phase (3) of the algorithm. Given the $(3(f + 1), 0)$ $\texttt{FT-ABFS}$ structure T_2, a subgraph $(3(f + 1), (f + 1) \log n))$ $\texttt{FT-ABFS}$ structure T_3 of $O(n)$ edges is constructed as follows. Let $G' = T_2 \setminus T_0$ be the subgraph obtained by removing all original BFS edges of T_0 from T_2. Let $T' \leftarrow \texttt{ConsSpan}(G', \log n, f)$ be an f-edge fault tolerant $(\log n, 0)$ spanner for G'. The resulting structure is $H = T_0 \cup T'$.

Analysis. We first provide an auxiliary claim regarding the labeling scheme $\texttt{Label}_F(u_i)$, $F \subseteq E$.

Observation 2.14 *(1) $\pi(u_i, u_{i'}) \subseteq T_0 \setminus F$ for every pair of vertices $u_i, u_{i'}$ such that $\texttt{Label}_F(u_i) = \texttt{Label}_F(u_{i'})$.*

(2) $|\bigcup\{\texttt{Label}_F(u') \mid u' \in V\}| = |\mathcal{C}(T_0 \setminus F)| \leq |F| + 1$.

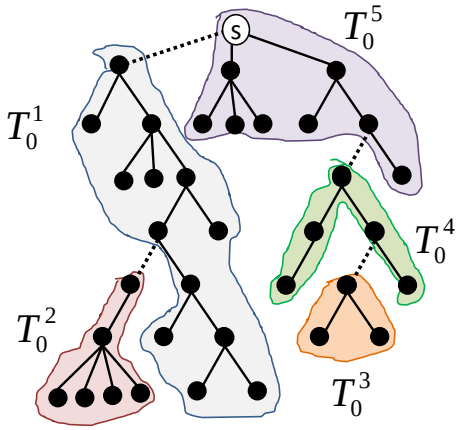

Figure 6: Upon the failure of the BFS edges F (dotted lines), the BFS tree T_0 is decomposed into $c = |F| + 1 \leq f + 1$ components, denoted $T_0^1, \ldots, T_0^c$.

Proof: Part (1) follows by definition, since two vertices are assigned the same label for a given edge fault F iff they belong to the same tree in the forest $T_0 \setminus F$. Part (2) is proven by induction. Let $e_1, \ldots, e_\ell \in F \cap E(T_0)$ be the faulty edges in T_0. We claim that $|\mathcal{C}(T_0 \setminus F)| = \ell + 1$. Since $\ell \leq |F|$, this would establish the observation. For the base of the induction, $\ell = 1$, note that every faulty edge $e = (x, y)$ disconnects the tree T_0 into two nonempty components, one containing x and one containing y, hence $|\mathcal{C}(T_0 \setminus \{e\})| = 2$. Assume this holds for every $\ell' \leq \ell$ and consider $\ell' + 1$. By the induction assumption, the number of components in $T' = T_0 \setminus \{e_1, \ldots, e_{\ell'}\}$ is $|\mathcal{C}(T')| = \ell' + 1$. Let $T'' \in \mathcal{C}(T')$ be the connected tree in T' that contains $e_{\ell'+1}$. By the definition of $\mathcal{C}(T')$, such T'' exists. Then by the induction base, $\mathcal{C}(T'' \setminus \{e_{\ell'+1}\}) = 2$, since T'' is broken into two components upon the removal of the edge $e_{\ell'+1}$. We get that $\mathcal{C}(T \setminus \{e_1, \ldots, e_{\ell'+1}\}) = |\mathcal{C}(T')| + 1 = \ell' + 2$. The observation follows. ∎

Fix a vertex u, edge faults $F \subseteq E$, $|F| \leq f$, and a replacement path $P = P^*_F(u)$. Let $\texttt{New}(P) = \{e_1, \ldots, e_\ell\}$, where $e_i = (u_i, v_i)$ for $i \in \{1, \ldots, \ell\}$ be the set of new edges in order of appearance on P (from s) and let $\texttt{New}^+(P) = \{e_{i_1}, \ldots, e_{i_{\ell'}}\} \subseteq \texttt{New}(P)$ be the corresponding ordered set taken into T_2. Let $\texttt{New}^-(P) = \texttt{New}(P) \setminus \texttt{New}^+(P)$ the set of new edges not included in T_2. The following observation is immediate by the structure of the algorithm.

Observation 2.15 *(1)* $e_{i_1} = e_1$.

(2) $M(v_{i_y}) \in V(P[v_{i_y}, v_{i_{y'}}]) \setminus \{v_{i_{y'}}\}$ *for every* $x' > x$.

(3) $P[M(v_{i_{y-1}}), v_{i_y}] \subseteq T_2 \setminus F$.

(4) $\texttt{Label}_F(v_{i_y}) \neq \texttt{Label}_f(v_{i_{y'}})$ *for every* $y, y' \in \{1, \ldots, \ell'\}$.

Proof: Parts (1) and (2) follow by the description of the algorithm. To see (3), note that e_{i_y} is the first new edge that appears after $M(v_{i_{y-1}})$, i.e., $e_{i_y} = \texttt{FirstNewE}(P[M(v_{i_{y-1}}), u])$. Hence, $\texttt{New}(P[M(v_{i_{y-1}}), v_{i_y}]) = \{e_{i_y}\}$. In addition, since $P[M(v_{i_{y-1}}), v_{i_y}]$ appeared on the replacement path $P \subseteq T_1 \setminus F$, it holds that $E(P[M(v_{i_{y-1}}), v_{i_y}]) \cap F = \emptyset$. Since e_{i_y} (the only new edge of $P[M(v_{i_{y-1}}), u]$ is taken to T_2, (3) follows. Finally, consider Part (4). Assume towards contradiction that there exists some $x < x'$ such that $\texttt{Label}_F(v_{i_y}) = \texttt{Label}_f(v_{i_{y'}})$. By Part (2), $M(v_{i_y}) \in V(P[v_{i_y}, v_{i_{y'}}]) \setminus \{v_{i_{y'}}\}$. Since $\texttt{Label}_F(v_{i_y}) = \texttt{Label}_f(v_{i_{y'}})$ and $v_{i_{y'}}$ appears *after* $M(v_{i_y})$, we get a contradiction to the selection of $M(v_{i_y})$. The observation follows. ∎

We proceed by showing that the multiplicative stretch introduced by excluding $\texttt{New}^-(P)$ from the intermediate structure T_2 is at most $3(f+1)$. For $P^*_F(u)$, define the corresponding $s-u$ replacement path $Q_f(u)$. Let $Q_1 = P[s, v_{i_1}] \circ \pi(v_{i_1}, M(v_{i_1}))$ and $Q_y = P[M(v_{i_{y-1}}), v_{i_y}] \circ \pi(v_{i_y}, M(v_{i_{y+1}}))$ for every $y \in \{2, \ldots, \ell' - 1\}$. Then

$$Q_f(u) = Q_1 \circ Q_2 \ldots \circ Q_{\ell'} \circ P[M(v_{i_{\ell'}}), u].$$

For pictorial illustration, see Fig. 7.

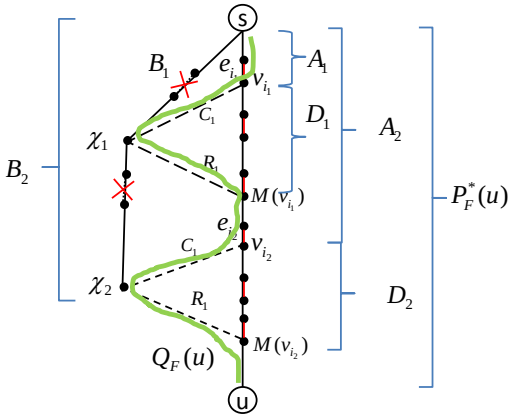

Figure 7: Illustration of the replacement paths used in T_2. Black lines correspond to BFS edges. Shown is the $P^*_F(u)$ replacement path where new edges $\texttt{New}(P^*_F(u))$ are marked in red. The dashed lines correspond to the BFS path used to bypass segments in P^*_F. The replacement $s-u$ path $Q_f(u)$ with at most $f+1$ new edges is marked in green.

We state that $Q_f(u)$ is an $s-u$ replacement path in $T_2 \setminus F$ and it is longer than the optimal replacement path $P^*_F(u)$ by a factor of at most $3(f+1)$.

Lemma 2.16 *(1)* $Q_F(u) \subseteq T_2 \setminus F$.

(2) $|Q_F(u)| \leq 3(f+1) \cdot |P^*_F(u)|$.

Proof: Let $Q^* = Q_f(u)$. We begin with (1). Note that the missing new edges $\texttt{New}(P)^-$ appear on P between some two vertices of the same label. Specifically, $\texttt{New}(P)^- \subseteq \bigcup_{y=1}^{\ell'} P[v_{i_y}, M(v_{i_y})]$. Since $\texttt{Label}_F(v_{i_y}) = \texttt{Label}_F(M(v_{i_y}))$, by Obs. 2.14, it holds that $\pi(v_{i_y}, M(v_{i_y})) \subseteq T_0 \setminus F$ and hence also $\pi(v_{i_y}, M(v_{i_y})) \subseteq T_2 \setminus F$. Formally, by Obs. 2.15(1) $e_{i_1} = e_1 = \texttt{FirstNewE}(P)$ hence $Q_1 \subseteq T_0 \setminus F$. By Obs. 2.15(3) the subpath $P[M(v_{i_{y-1}}), v_{i_y}] \in T_2 \setminus F$ for every $y \in \{2, \ldots, \ell' - 1\}$. Finally, by Cl. 2.14, $\pi(v_{i_y}, M(v_{i_{y+1}})) \subseteq T_0 \setminus F$ for every $y \in \{1, \ldots, \ell' - 1\}$, since $\texttt{Label}_F(v_{i_y}) = \texttt{Label}_F(M(v_{i_{y+1}}))$. It follows that $Q^* \subseteq T_2 \setminus F$.

We now turn to (2) and show that $|Q^*| \leq 3(f+1)|P|$. For every $y \in \{1, \ldots, \ell'\}$, we define the following in order to be able to bound the length of the bypass. Let $\chi_y = \texttt{LCA}(v_{i_y}, M(v_{i_y}))$, $A_y = P[s, v_{i_y}], B_y = \pi(s, \chi_y), C_y = \pi(\chi_y, v_{i_y}), D_y = P[v_{i_y}, M(v_{i_y})]$ and $R_y = \pi(\chi_y, M(v_{i_y}))$. Let $\widetilde{P}_y = A_y \circ C_y \circ R_y$ and $P_y = P[s, M(v_{i_y})]$.

Claim 2.17 *$|\widetilde{P}_y| \leq 3 \cdot |P_y|$ for every $y \in \{1, \ldots, \ell'\}$.*

Proof: Note that since $B_y \circ C_y = \pi(s, v_{i_y})$, the alternative $s - v_{i_y}$ path A_y satisfies $|A_y| \geq |B_y| + |C_y|$. In addition, since $C_y \circ D_y$ is an alternative $\chi_y - M(v_{i_y})$ path and $R_y = \pi(\chi_y, M(v_{i_y}))$ is the $\chi_y - M(v_{i_y})$ shortest-path, it holds that $|R_y| \leq |C_y| + |D_y|$. Hence, $|\widetilde{P}_y| \leq |A_y| + 2|C_y| + |D_y| \leq 3|A_y| + |D_y| \leq 3|P_y|$. ∎

We now claim that the replacement $s - M(v_{i_y})$ path

$$\widetilde{Q}_y = Q_1 \circ Q_2 \ldots \circ Q_y \subseteq T_2 \setminus F$$

satisfies $|\widetilde{Q}_y| \leq 3y|P_y|$. This is proved by induction on y. For the base of the induction consider $y = 1$. In this case $\widetilde{Q}_1 = \widetilde{P}_1$. Hence, the claim follows by Cl. 2.17 and $|\widetilde{Q}_1| \leq 3 \cdot |P_1|$. Assume the claim holds up to $y - 1$ and consider y.

$$\begin{aligned}
|\widetilde{Q}_y| &= |\widetilde{Q}_{y-1}| + |Q_y| \leq 3(y-1)|P_{y-1}| + |P[M(v_{i_{y-1}}), v_{i_y}]| + |C_y \circ R_y| \\
&\leq 3(y-1)|P[s, v_{i_y}]| + |C_y \circ R_y| \\
&= 3(y-1)|A_y| + |C_y \circ R_y| \leq 3(y-1)|A_y| + 3|A_y \circ D_y| \\
&\leq 3y \cdot |P_y| \ ,
\end{aligned}$$

where the first inequality follows by the induction assumption, as $|\widetilde{Q}_{y-1}| \leq 3(y-1)|P_{y-1}|$, and the third inequality follows by Cl. 2.17. Since by Obs. 2.15(4) and Obs. 2.14(2), $\ell' \leq f + 1$, we get that $|\widetilde{Q}_{\ell'}| \leq 3 \cdot (f+1)|P_{\ell'}|$ and $|Q^*| = |\widetilde{Q}_{\ell'}| + |P[M(v_{i_{\ell'}}), u]| \leq 3(f+1)|P_{\ell'}| + |P[M(v_{i_{\ell'}}), u]| \leq 3(f+1)|P|$. The lemma follows. ∎

We therefore have the following.

Corollary 2.18 *For every u and $F \subseteq E$, there exists a replacement path $Q_F(u) \in T_2 \setminus F$ such that*

(1) $|\texttt{New}(Q_F(u))| \leq f+1$ *and*

(2) $|Q_F(u)| \leq 3(f+1) \cdot P^*_F(u)$.

Hence T_2 *is an* f*-edge* $(3(f+1),0)$ `FT-ABFS` *structure.*

Finally, we prove the correctness of the last phase and show the following.

Claim 2.19 *(1)* H *is an* f*-edge* $(3(f+1),(f+1)\log n)$ `FT-ABFS` *structure with respect to* s.

(2) $|H| = O(fn)$.

Proof: Let $P = P^*_F(u)$ be the optimal replacement path and let $Q = Q_F(u)$ be the corresponding replacement path in T_2 obtained by using at most $f+1$ bypasses between vertices of the same label on P. Then, by Cor. 2.18, $|\texttt{New}(Q)| \leq f+1$ and $|Q| \leq 3(f+1) \cdot P$. Let $G' = T_2 \setminus T_0$ and $G'' = \texttt{ConsSpan}(G', \log n, f)$ be the f-edge FT $(\log n, 0)$ spanner for G' (see Fact 2.13). Let $E' = \texttt{New}(Q) \setminus H$ be the set of new edges in Q that are missing in the final H. Hence, $E' \subseteq G' \setminus G''$. Since $E' \subseteq \texttt{New}^+(Q)$, it holds that $|E'| \leq |\texttt{New}^+(Q)| \leq f+1$. We now claim that for every missing edge $e = (x,y) \in E'$ there exists an $x-y$ path in the surviving structure $G'' \setminus F$ of length at most $\log n$. By the fact that G'' is an f-edge $(\log n, 0)$ spanner for G', it holds that $\text{dist}(x,y,G'' \setminus F) \leq \log n \cdot \text{dist}(x,y,G' \setminus F) = \log n$, where the last equality follows by the fact that $e = (x,y)$ appears on the replacement path Q, hence $e \notin F$, so $\text{dist}(x,y,G' \setminus F) = 1$. We therefore have that Q contains at most $f+1$ missing edges, and for each there exists a path in G" of length at most $\log n$. Overall, $\text{dist}(s,u,G \setminus F) \leq |E(Q) \setminus E'| + (\log n) \cdot |E'| \leq 3(f+1)|P| + (f+1)\log n$. Part (1) is established. We now consider part (2). By Fact 2.13, $|G''| = O(n)$, hence $H = T_0 \cup G''$ has $O(fn)$ edges as well. The claim follows. ∎

This completes the proof of Thm. 2.12(1). We now consider part (2) of the theorem asserting that it is possible to get rid of the additive factor, albeit at the expense of considerably increasing the size of the `FT-ABFS` structure.

Proof: [Thm. 2.12(2)] For a given $k \geq 3$, we first construct the collection of all replacement paths $P^*_F(u_i) \in SP(s,u_i,G \setminus F, W)$ for every $F \subseteq E$ and every $u_i \in V$. Let $\ell = (f+1) \cdot (2k-1)$. The $(3(f+1)+1, 0)$ `FT-ABFS` structure H is constructed in two steps. A replacement path $P^*_F(u)$ is *short* iff $|P^*_F(u)| \leq \ell$. In the first step, a subgraph $H_1 \subseteq G$ is constructed containing the set of all short replacement paths, i.e., $H_1 = \{\texttt{LastE}(P^*_F(u_i)) \mid |P^*_F(u_i)| \leq \ell\}$. In the second step, an $(3(f+1),\ell)$ `FT-ABFS` structure H_2 is constructed by employing the algorithm of Thm. 2.12(1) with one minor modification; in step (3) of the algorithm we construct an $(2k-1,0)$ spanner instead of $(\log n, 0)$ spanner, i.e., step (3) is given by $T' \leftarrow \texttt{ConsSpan}(T_2 \setminus T_0, 2k-1, f)$ for the given k (while $2k-1 = \log n$ in the original algorithm). By Fact 2.13 and the proof of

Cl. 2.19, it holds that $|E(H_2)| = O(f \cdot n + n^{1+1/k})$. We next bound the size of H_1.

Claim 2.20 $|H_1| \leq n \cdot ((f+1) \cdot (2k-1))^{f+1}$.

Proof: Let $\mathcal{P}_i = \{P^*_F(u_i), F \subseteq E\}$ be the collection of $s - u_i$ replacement paths. Define $\mathcal{P}_i^0 = \{\pi(s, u_i)\}$ and $\mathcal{P}_i^{f'} = \{P^*_F(u_i) \mid |F| = f'\}$ for $f' \geq 1$ as the collection of $s - u_i$ replacement paths supporting a sequence of f' edge faults. Hence, $\mathcal{P}_i = \bigcup_{f'=1}^{f} \mathcal{P}_i^{f'}$.

We prove that for every $f' \in \{1, \ldots, f\}$ if holds that $|\mathcal{P}_i^{f'}| \leq \ell \cdot |\mathcal{P}_i^{f'-1}|$. To see this, observe that every replacement path $P^*_F(u_i) \in \mathcal{P}_i^{f'} \setminus \mathcal{P}_i^{f'-1}$ protects an edge failure e in some $P^*_{F'}(u_i) \in \mathcal{P}_i^{f'-1}$. Hence, $F = F' \cup \{e\}$. Since $\mathcal{P}_i^{f'-1}$ contains only short replacement paths, it holds that each short $s - u_i$ replacement path $P^*_{F'}(u_i) \in \mathcal{P}_i^{f'-1}$ has at most ℓ replacement paths of the form $P^*_{F' \cup \{e\}}(u_i)$ in $\mathcal{P}_i^{f'}$ that protect against the failure of $e \in P^*_{F'}(u_i)$. Concluding that $|\mathcal{P}_i^{f'}| \leq \ell \cdot |\mathcal{P}_i^{f'-1}|$. Overall, $|\texttt{New}_i| \leq |\bigcup_{f'=1}^{f} \mathcal{P}_i^{f'}| \leq \ell^{f+1}$ and $|H_1| \leq n \cdot \ell^{f+1}$. The claim follows. ∎

Finally, we show that $H = H_1 \cup H_2$ is an $(3(f+1)+1, 0)$ FT-ABFS structure. Consider a vertex edge-set pair (i, F) corresponding to a vertex $u_i \in V$ and edge set $F \subseteq E$. There are two cases. Case (a) is where the replacement path $P^*_F(u_i)$ is short, i.e., $|P^*_F(u_i)| \leq \ell$. By the same argumentation as in Lemma 2.6 it holds that $\text{dist}(s, u_i, H_1 \setminus F) = \text{dist}(s, u_i, G \setminus F)$ (since H_1 contains the last edges of these replacement paths, which was shown to be sufficient). The complementary case is where the replacement path is long, $|P^*_F(u_i))| > \ell$. By the proof of Thm. 2.12(1), it holds that $\text{dist}(s, u_i, H_2 \setminus F) \leq (3(f+1))\text{dist}(s, u_i, G \setminus F) + \ell \leq (3(f+1)+1)\text{dist}(s, u_i, G \setminus F)$. The claim follows. ∎

3 Lower Bounds for Additive FT-ABFS Structures

In this section we provide lower bound constructions for $(1, \beta)$ FT-ABFS structures for various values of the additive stretch β. The starting point for these constructions is the lower bound construction for the exact FT-BFS structures in [15]. In this construction, the bulk of the edges was due to a complete bipartite graph $B = (L, R)$ where $|L| = \Theta(n)$ and $|R| = \Theta(\sqrt{n})$. When relaxing to additive stretch 2, it is no longer required to include B entirely in the spanner; in fact, one can show that it suffices to include in the $(1, 2)$ FT-ABFS structure only a subgraph $B' \subseteq B$ of $O(n)$ edges. This subgraph B' is obtained by connecting an arbitrary vertex $\ell^* \in L$ to every vertex $r \in R$ and connecting an arbitrary vertex $r^* \in R$ to every vertex $\ell \in L$. Thus our goal is to replace B with a dense subgraph of high girth. Constructions of this type are known, but our requirement is in fact more stringent. Note that an essential characteristic of the construction of [15] is that the R layer consists of at most $c \cdot \sqrt{n}$ vertices, for some constant $c > 0$. A larger layer could not be supported since the upper bound analysis of [15] implies that

every vertex in L requires at most $O(\sqrt{n})$ edges in every FT-BFS structure. Since the known lower bound on the number of edges in a $\Theta(n) \times \Theta(\sqrt{n})$ bipartite graph with high girth (greater than 4) is $O(n)$, such graphs are not good candidates for replacing the bipartite graph B in the construction, as using them results in graphs of $O(n)$ rather than $O(n^{1+\epsilon(\beta)})$ edges. Instead, our constructions replace the complete *unbalanced* bipartite graph B by a *multiple* copies of *balanced* $\Theta(\sqrt{n}) \times \Theta(\sqrt{n})$ bipartite graphs with high girth. Hence, the lower bound constructions of $(1, \beta)$ FT-ABFS structure rely on known construction of balance bipartite graphs with high girth. The desired construction is achieved by carefully inserting multiple copies of these graphs.

Additive Stretch $\beta = 1$. By a proof very similar to that of the exact case (see [15]) one can show the following (details are omitted).

Theorem 3.1 *There exists an n-vertex graph $G(V, E)$ and a source $s \in V$ such that any $(1, 1)$ FT-ABFS structure rooted at s has $\Omega(n^{3/2})$ edges.*

Additive stretch $3 \leq \beta \leq O(\log n)$. The *girth* of a graph G is the minimum number of edges on any cycle in G. Let $B(n, g) = (L, R, E)$ be a bipartite graph, where $|L| = |R| = n$, the girth is at least g and the number of edges is $|E| = \Omega(n^{m(g)})$ for some function $m(g) > 1$. Removing any edge (u, v) in such a graph $B(n, g)$ increases the distance between its endpoints u and v from 1 to $g - 1$, which implies that any (strict) subgraph of $B(n, g)$ has additive stretch $\beta \geq g - 2$.

In what follows, we embed copies of $B(d, g)$ graphs for some parameters d and g in order to achieve a lower bound for the case of a $(1, \beta)$ FT-ABFS structure with *edge faults*. We show the following.

Theorem 3.2 *For every integer n and constant integer β, there exists an n-vertex graph $G(n, \beta) = (V, E)$ and a source $s \in V$ such that any $(1, \beta)$ FT-ABFS structure with respect to s, H, has:*

(1) $|E(H)| = \Omega(n^{5/4})$, *if* $\beta \leq 3$.

(2) $|E(H)| = \Omega(n^{1+1/2(\beta+3)})$, *if* $4 \leq \beta \leq O(\log n)$.

(3) $|E(H)| = \Omega(n^{7/6})$, *if* $\beta \leq 5$.

(4) $|E(H)| = \Omega(n^{11/10})$, *if* $\beta \leq 9$.

Proof: Given integers n and β, the graph $G(n, \beta) = G(V, E)$ consists of the following main components (illustrated in Fig. 8):

(1) A path $P_0 = [s = v_1, \ldots, v_{d+1} = v^*]$ of length d, for an integer parameter d fixed later. Our focus in the analysis is on what happens when some edge on this path fails.

(2) A set $X = \bigcup_{i=1}^{d} X_i$ of d^2 vertices organized in d sets $X_1, \ldots, X_d$ each of size d, where $X_i = \{x_{i,1}, \ldots, x_{i,d}\}$ for every $i \in \{1, \ldots, d\}$.

(3) A set of d^2 vertex disjoint paths of length $\beta+1$ each, $U_{i,j} = [v^* = u^i_1, \ldots, u^i_{\beta+2} = x_{i,j}]$, for every $i, j \in \{1, \ldots, d\}$ connecting the vertex $v^* \in P_0$ to the vertices of X.

(4) A set $Z = \bigcup_{i=1}^{d} Z_i$ of d^2 vertices organized in d sets $Z_1, \ldots, Z_d$ each of size d, where $Z_i = \{z_{i,1}, \ldots, z_{i,d}\}$ for every $i \in \{1, \ldots, d\}$.

(5) A vertex set $W = \{w_1, \ldots, w_d\}$, where vertex w_j is connected by d vertex disjoint paths $Q_{i,j}$ of length-$\beta + 1$ to the d vertices $\{z_{i,j} \in Z_i \mid i = 1, \ldots, d\}$. Altogether there exist d^2 vertex disjoint paths $Q_{i,j} = [w_j = q^j_1, \ldots, q^j_{\beta+2} = z_{i,j}]$, one for every $i, j \in \{1, \ldots, d\}$.

(6) A collection of d vertex disjoint paths of decreasing length, $P_1, \ldots, P_d$, where for $j \in \{1, \ldots, d\}$, $P_j = [v_j = p^j_1, \ldots, p^j_{\ell_j} = w_j]$ connects v_j with w_j and its length is $\ell_j = |P_j| = d + 4 + (\beta + 1) \cdot (d - j + 1)$.

(7) d copies of the bipartite graph $B(d, \beta+3)$, where $B_i(d, \beta+3) = (X_i, Z_i, E_i)$ connects X_i to Z_i for every $i \in \{1, \ldots, d\}$. This graphs contribute most of the edges in G (with all the other components containing only $O(n)$ edges).

(8) A path $R = [s = r_1, \ldots, r_{d'+1}]$ for $d' \geq 0$ to be defined later, added in order to complete the number of vertices in G to exactly n. (This path has no special role in the construction.)

Overall, the vertex set of $G(n, \beta)$ is

$$V = X \cup Z \cup V(P_0) \cup \left(\bigcup_{i=1}^{d} V(P_i) \right) \cup \bigcup_{i=1}^{d} \bigcup_{j=1}^{d} (V(Q_{i,j}) \cup V(U_{i,j})) \cup V(R).$$

and its edge set is

$$E = \bigcup_{i=1}^{d} E(B_i(d, \beta + 3)) \cup E(P_0) \cup E(R) \cup \bigcup_{i=1}^{d} E(P_i) \cup \bigcup_{i,j} (E(Q_{i,j}) \cup E(U_{i,j})).$$

Set $d = \lfloor \sqrt{n/(14\beta)} \rfloor$, and let $d' = n - |V \setminus (R \setminus \{s\})|$. We now turn to prove the correctness of the construction and establish Thm. 3.2. Note that without the path R, the rest of $G(n, \beta)$ contains fewer than n vertices, as $|V(P_0)| = d$, $|X \cup Z| = 2d^2$,

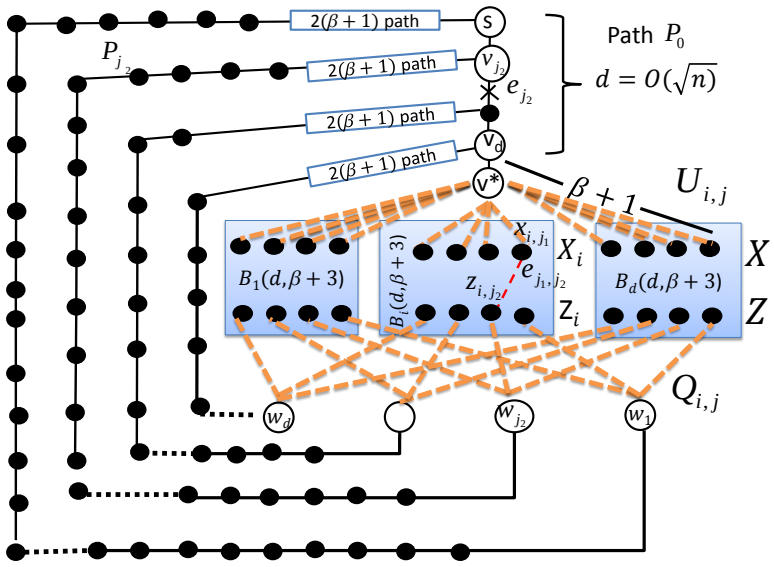

Figure 8: Lower bound construction for FT-ABFS structures. The original BFS tree consists of the non-dashed edges. The dotted edges are necessary to make it an FT-ABFS structure. The orange dashed lines correspond to $\beta+1$ length path. The edges of the bipartite graphs $B_i(d,\beta+3)$ are required in every $(1,\beta)$ FT-ABFS structure with respect to s. For example the red edge e_{j_1,j_2} is necessary upon the fault of $e_{j_2} \in P_0$.

$|\bigcup_{i=1}^{d} V(P_i)| = \sum_{j=1}^{d} (d+4+(\beta+1)(d-j+1)) \leq 10\beta \cdot d^2$, $|\bigcup_{i,j} V(U_i)| = d^2 \cdot (\beta+2)$ and $|\bigcup_{i,j} V(Q_{i,j})| = d^2 \cdot (\beta+2)$. Hence overall, $|V \setminus (R \setminus \{s\})| \leq 14\beta \cdot d^2 \leq n$. Fix $d' = |R|$ to complete the size of $G(n,\beta)$ to n, i.e., set $d' = n - |V \setminus (R \setminus \{s\})|$.

Observation 3.3 $|E| = \Omega(n^{1/2+m(\beta+3)/2})$.

Proof: The bulk of the edges is due to the bipartite graphs $B_i(d,\beta+3)$. Since there are d such graphs, each with $\Omega(d^{m(\beta+3)})$ edges, we get that $|E| \geq \sum_i E(B_i(d,\beta+3)) = d \cdot \Omega(d^{m(\beta+3)}) = \Omega(n^{1/2+m(\beta+3)/2})$. The Claim follows. ∎

We now show the following.

Claim 3.4 *Every* $(1,\beta)$ FT-ABFS *structure* H *for* G *with respect to* s*, must contain all the edges of* $B_i(d,\beta+3)$ *for every* $i \in \{1,\ldots,d\}$.

Proof: Let $G = G(n,\beta)$. Assume, towards contradiction, that there exists some $i \in \{1,\ldots,d\}$ and an $(1,\beta)$ FT-ABFS structure H for G such that $B_i(d,\beta+3) \nsubseteq H$. Let $e_j = (v_j, v_{j+1})$ for every $j \in \{1,\ldots,d\}$. Let $e_{j_1,j_2} = (x_{i,j_1}, z_{i,j_2})$, where $x_{i,j_1} \in X_i$ and $z_{i,j_2} \in Z_i$, be a missing edge in H. Note that for $i',j' \in \{1,\ldots,d\}$, the shortest $s - z_{i',j'}$ path in $G \setminus \{e_{j'}\}$ is $P_{i',j'} = P_0[s = v_1, v_{j'}] \circ P_{j'} \circ Q_{i',j'}$.

Upon the failure of the edge $e_{j_2} = (v_{j_2}, v_{j_2+1}) \in P_0$, the unique shortest $s - x_{i,j_1}$ path in $G \setminus \{e_{j_2}\}$ is $P^* = P_{i,j_2} \circ e_{j_1,j_2}$. Since $e_{j_1,j_2} \notin H$ and therefore also $P^* \nsubseteq H$, the $s - x_{i,j_1}$ distance in $H \setminus \{e_{j_2}\}$ is strictly larger than that in $G \setminus \{e_{j_2}\}$. Let $P' \in SP(s, x_{i,j_1}, H \setminus \{e_{j_2}\})$ be the $s - x_{i,j_1}$ shortest-path in $H \setminus \{e_{j_2}\}$. By construction, P' must traverse the Z vertices (as the path P_0 is disconnected).

Let $z_{i',j'}$ be the first Z vertex occurring on P'. There are three cases to consider.
Case (C1) $j' > j_2$ and $i' \in \{1, \ldots, d\}$. Since the faulty edge disconnects the shortest-path of G between s and $v_{j'}$ for every $j' > j_2$, the shortest-path between s and $z_{i',j'}$ in $G \setminus \{e_{j_2}\}$ must visit another Z vertex, in contradiction to the fact that $z_{i',j'}$ is the first Z vertex on P'.
Case (C2) $j' < j_2$ and $i' \in \{1, \ldots, d\}$. In this case, the shortest-path between s and $z_{i',j'}$ in $G \setminus \{e_{j'}\}$ is $P_{i',j'}$. Since $|P_{j'}| \geq |P_{j_2}| + \beta + 1$ for every $j' < j_2$, it holds that $|P'| > |P^*| + \beta + 1$, in contradiction to the fact that H is a $(1, \beta)$ `FT-ABFS` structure.
Case (C3) $j' = j_2$. This case is further divided into two cases.
Subcase(C3A): $i \neq i'$. The shortest-path P' must be of the form $P' = P_{i',j_2} \circ Q'$ where $Q' \in SP(z_{i',j_2}, x_{i,j_1}, H \setminus \{e_{j_2}\})$. Since any path connecting vertex $w \in B_i(d, \beta + 3)$ and $w' \in B_{i'}(d, \beta + 3)$ is of length at least $\beta + 1$, it holds that $|P'| \geq |P^*| + \beta + 1$, and we end with contradiction again.
Subcase (C3B): $i = i'$. Since the edge $e_{j_1,j_2} \notin T'$ is missing, z_{i,j_2} is the first Z vertex on P', and the distance between $x_{i,j} \in B_i(d, \beta + 3)$ and $w' \in B_{i'}(d, \beta + 3)$ for $i \neq i'$, is of length at least $\beta + 1$, it follows that x_{i,j_2} is connected to z_{i,j_2} via vertices in $B_i(d, \beta + 3)$, i.e., $P' = P_{i,j_2} \circ Q'$ where Q' is an $z_{i,j_2} - x_{i,j_1}$ shortest path in H. Since the girth of $B_i(d, \beta + 3)$ is $g \geq \beta + 3$, it holds that $\mathrm{dist}(x_{i,j_1}, z_{i,j_2}, B_i(d, \beta + 3) \setminus \{e_{j_1,j_2}\}) \geq \beta + 2$, hence $|Q'| \geq \beta + 2$ and $|P'| \geq |P^*| + \beta + 1$, contradiction. ∎

Corollary 3.5 *Every $(1, \beta)$ `FT-ABFS` structure H for G must contain at least $d \cdot d^{m(\beta+3)}$ edges.*

To establish the theorem we now make use of the following fact concerning the existence of graphs with "many" edges and high girth.

Fact 3.6 *For every sufficiently large n there exists a connected bipartite graph $B(n, g) = \{L, R, E\}$, $|L| = |R| = n$ such that*

(1) $m(g) = 3/2$ for $g = 6$, Lemma 5.1.1 of [13],

(2) $m(g) = 4/3$ for $g \geq 8$ [14]

(3) $m(g) = 6/5$ for $g \geq 12$ [14] and

(4) $m(g) = 1 + 1/g$ for every $6 \leq g \leq 2 \log n$, [17].

Thm. 3.2 follows by Cor. 3.5 and by applying Fact 3.6. ∎

4 Upper Bound for Additive Stretch 4

In this section, we establish the following.

Theorem 4.1 *There exists a poly-time algorithm that for every n-vertex unweighted undirected graph G and source s constructs a $(1,4)$* FT-ABFS *structure H with $O(n^{4/3})$ edges.*

As usual, the starting point of our construction is the shortest path tree T_0, and this tree should be fortified against possible failure of any of its edges, by adding edges and augmenting T_0 into H.

Overview. The construction of H consists of 3 main stages and revolves around the following key observation. Let $E' \subseteq E(G)$ be a subset of edges. A replacement path $P^*_{i,j} \in SP(s, v_i, G \setminus \{e_j\})$ is "missing-ending" with respect to E' if the last edge of $P^*_{i,j}$ does not belong to E', i.e., $\texttt{LastE}(P^*_{i,j}) \notin E'$. Let $\mathcal{P}^* = \{P^*_{i,j} \mid v_i \in V \text{ and } e_j \in \pi(s, v_i)\}$ be the collection of all $s - v_i$ replacement paths for every $e_j \in \pi(s, v_i)$ and consider its partition $\mathcal{P}^*_{E'} = \mathcal{P}^{miss}_{E'} \cup \mathcal{P}^+_{E'}$ into missing-ending paths $\mathcal{P}^{miss}_{E'} = \{P^*_{i,j} \mid \texttt{LastE}(P^*_{i,j}) \notin E'\}$ and non missing-ending paths $\mathcal{P}^+_{E'} = \{P^*_{i,j} \mid \texttt{LastE}(P^*_{i,j}) \in E'\}$. Clearly, a subgraph H containing all replacement paths $\mathcal{P}^*$ is an exact FT-BFS structure, however by [15] such a subgraph might contain $\Omega(n^{3/2})$ edges. One of the key observations in our analysis is that it is sufficient to "take care" only of the missing ending path collection $\mathcal{P}^{miss}_{E'}$.

Observation 4.2 *Let $H \subseteq G$ be a subgraph containing E' and satisfying $\text{dist}(s, v_i, H \setminus \{e_j\}) \leq \text{dist}(s, v_i, G \setminus \{e_j\}) + \beta$ for every $v_i \in V$ and $e_j \in \pi(s, v_i)$ such that $P^*_{i,j} \in \mathcal{P}^{miss}_{E'}$. Then H is a $(1, \beta)$* FT-ABFS *structure, i.e., $\text{dist}(s, v_i, H \setminus \{e_j\}) \leq \text{dist}(s, v_i, G \setminus \{e_j\}) + \beta$ also holds for non missing-ending paths $P^*_{i,j} \in \mathcal{P}^+_{E'}$.*

Observation 4.2 provides the basis for the general structure of the algorithm. Its proof is based on ideas resembling those appearing in the proof of Lemma 4.19).

A vertex-edge pair (i, j), representing a vertex v_i and an edge $e_j \in \pi(s, v_i)$, is *satisfied* by a subgraph $H \subseteq G$ if there exists a replacement path $P \in SP(s, v_i, G \setminus \{e_j\})$ whose last edge $\texttt{LastE}(P)$ is in H, otherwise it is *unsatisfied*. For a subgraph $H \subseteq G$, let

$$\begin{aligned}\Gamma(H) = \{&(i,j) \mid e_j \in \pi(s, v_i) \text{ and} \\ &\texttt{LastE}(P) \notin H \text{ for every } P \in SP(s, v_i, G \setminus \{e_j\})\}\end{aligned}$$

be the collection of unsatisfied pairs in H. Starting with the collection of all pairs $\Gamma_0 = \{(i,j) \mid e_j \in \pi(s,v_i)\}$ that are required to be satisfied in H, the algorithm consists of three stages, aiming towards increasing the set of satisfied pairs by adding a suitable collection of edges to the constructed FT-ABFS structure H. Specifically, in each stage k, the algorithm is given a "partial" FT-ABFS structure $\mathcal{F}_k$ and a list of pairs $\Gamma_k \subseteq \Gamma_{k-1}$ that might not be satisfied yet in $\mathcal{F}_k$. Essentially, the pairs $\Gamma_{k-1} \setminus \Gamma_k$ are satisfied in $\mathcal{F}_k$. The algorithm then defines a subset of target pairs $\Delta_k \subseteq \Gamma_k$ and a corresponding collection of edges $E(\Delta_k)$ that aim to satisfy Δ_k in the final spanner. At the end of this stage, the algorithm sets $\mathcal{F}_{k+1} = \mathcal{F}_k \cup E(\Delta_k)$ and the updated list of unsatisfied pairs is reduced to $\Gamma_{k+1} = \Gamma_k \setminus \Delta_k$. To compute a sparse $E(\Delta_k)$, the algorithm considers for every pair $(i,j) \in \Delta_k$, a *specific* replacement path $\widetilde{P}_{i,j} \in SP(s,v_i,G \setminus \{e_j\})$. Only the last edge $\texttt{LastE}(\widetilde{P}_{i,j})$ is added to $E(\Delta_k)$. Hence, $E(\Delta_k) = \{\texttt{LastE}(\widetilde{P}_{i,j}) \mid (i,j) \in \Delta_k\}$. Finally, in the last stage 3, the replacement paths of the yet unsatisfied pairs in the current FT-ABFS structure $\mathcal{F}_3$ are considered to be added *entirely* to H, by employing a modified path-buying procedure. We now describe these stages in more detail. For a schematic illustration of the scheme, see Fig. 9.

At the first stage S1 of Subsec. 4.1, the algorithm clusters some of the vertices of $V(G)$, resulting in a clustered graph $G_C \subseteq G \setminus T_0$ which is shown to have $O(n^{4/3})$ edges and a set $\mathcal{C}$ of $O(n^{2/3})$ clusters. For this we use the clustering algorithm of [10], where not all the vertices of $V(G)$ are clustered but the only G edges missing in G_C are those incident to clustered vertices. Hence, letting $\Delta_1 = \{(i,j) \mid v_i \text{ is not clustered and } e_j \in \pi(s,v_i)\}$, the edge set $E(\Delta_1) = G_C$ satisfies Δ_1 in the current spanner $\mathcal{F}_1$. It therefore follows that it remains to handle only pairs (i',j) for clustered vertices $v_{i'}$. At the end of stage S1 (Sec. 4.2), the algorithm uses the clustering to divide the shortest path $\pi(s,v_i)$ of every clustered vertex v_i into three consecutive segments

$$\pi(s,v_i) = \pi^{far}(v_i) \circ \pi^{mid}(v_i) \circ \pi^{near}(v_i),$$

where the breakpoints depend on the clustering. Let $C(v_i) \in \mathcal{C}$ be the cluster of v_i and $\texttt{LCA}(C(v_i))$ be the least common ancestor in T_0 of the members of $C(v_i)$. Then, our segmentation satisfies that $\pi^{far}(v_i) \circ \pi^{mid}(v_i) = \pi(s, \texttt{LCA}(C(v_i)))$ and that $|\pi^{mid}(v_i)| \leq \lceil n^{2/3} \rceil$, and the vertices of $\pi^{far}(v_i)$ are at distance at least $n^{2/3}$ from v_i.

Equipped with this segmentation, the algorithm now handles separately edge failures $e_j \in T_0$ in each of these three segments. Specifically, stage S2 (in Sec. 4.3) deals with failures of $e_j \in \pi^{far}(v_i) \cup \pi^{near}(v_i)$, and stage S3 (in Sec. 4.4) deals with failures of $e_j \in \pi^{mid}(v_i)$ by employing a modified path-buying procedure.

Generally speaking, the edge faults in the far and near segments are handled by adding the collection of all last edges of the corresponding replacement paths. For a suitable construction of the replacement paths, this last edge collection is shown to be small. In contrast, the edge faults in the mid segments are handled by considering every

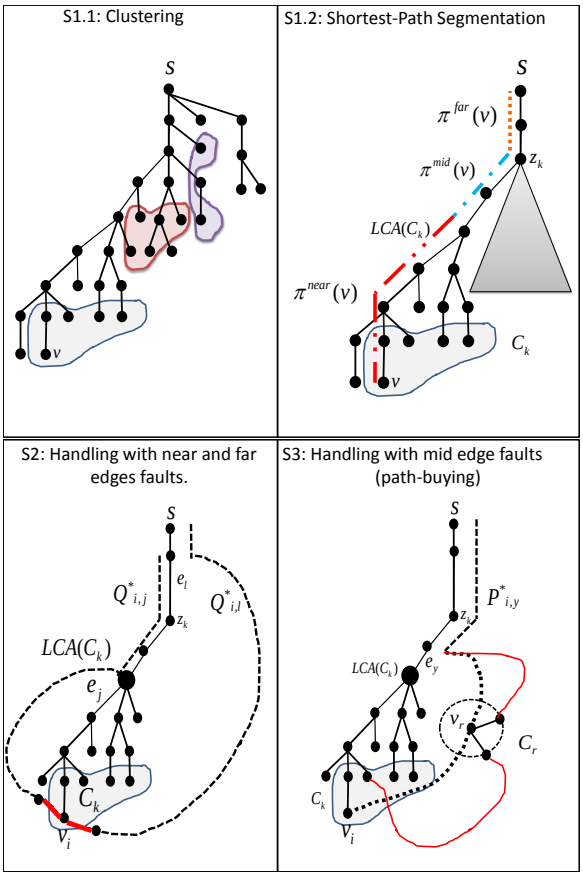
S1.1: Clustering
S1.2: Shortest-Path Segmentation
S2: Handling with near and far edges faults.
S3: Handling with mid edge faults (path-buying)

Figure 9: Schematic illustration of the main stages in H construction. Stage S1.1: Clustering in the graph $G \setminus T_0$. Stage S1.2: Segmentation of the shortest-path $\pi(s, v_i)$ into three segments: near, mid and far. Stage S2: Handling with the near and far edge faults by adding the last edge of their corresponding replacement paths $Q^*_{i,j}$. Stage S3: Handling with mid edge faults by applying a path buying procedure on the detour segments $P^-_{i,j}$ of the corresponding replacement paths $P^*_{i,j}$. A candidate detour is not added to H if there exists an alternative "safe" and short path between some clusters representatives.

replacement path $P^*_{i,j} \in SP(s, v_i, G \setminus \{e_j\})$ satisfying that $e_j \in \pi^{mid}(v_i)$ and adding it *entirely* to H if it satisfies some *cost* to *value* balance. To efficiently handle the faults, either in the near and far sections or in the mid section, a (nontrivial) preprocessing step of replacement path construction is required.

Let $IC = \{i \mid v_i \text{ is clustered}\}$. In stage S2, the algorithm defines, for every $i \in IC$, $\Delta^{near}_i = \{(i,j) \mid e_j \in \pi^{near}(v_i)\}$ and $\Delta^{far}_i = \{(i,j) \mid e_j \in \pi^{far}(v_i)\}$. The goal of this stage is to satisfy the pairs of $\Delta^{near} = \bigcup_{i \in IC} \Delta^{near}_i$ and $\Delta^{far} = \bigcup_{i \in IC} \Delta^{far}_i$ in the constructed `FT-ABFS` structure H. This stage consists of two substages.

In Substage S2.1, a collection of replacement path $Q^*_{i,j}$ is constructed. In Substage S2.2, the algorithm creates a sparse set $E(\Delta^{near})$ (resp., $E(\Delta^{far})$) containing the last edges of some replacement paths $Q^*_{i,j} \in SP(s, v_i, G \setminus \{e_j\})$ for every $(i,j) \in \Delta^{near}$ (resp., $(i,j) \in \Delta^{far}$). Due to some nice properties of the $Q^*_{i,j}$ collection, the analysis shows that $\left|E(\Delta^{near}) \cup E(\Delta^{far})\right| = O(n^{4/3})$. In particular, the size of the set $E(\Delta^{near}_i)$, consisting of the new edges appearing as last edges on $Q^*_{i,j} \in SP(s, v_i, G \setminus \{e_j\})$ satisfying that $e_j \in \pi^{near}(v_i)$, is bounded by a constant (which will be shown to follow from Fact. 4.3 that diameter of the clusters is constant). In addition, the number of last edges $E(\Delta^{far}_i)$ appearing on paths $\widetilde{P}_{i,j} \in SP(s, v_i, G \setminus \{e_j\})$ that protect against failures e_j in the far segment $\pi^{far}(v_i)$, is bounded by $O(n^{1/3})$ for every clustered vertex v_i, and as there are n vertices, in total there are $O(n^{4/3})$ edges in $E(\Delta^{far})$. This holds since one can show that the detours of the replacement paths for every clustered vertex v_i are *long*, *vertex disjoint*, and fully contained in the graph G. Hence, the remaining pairs (i,j) that should be handled in the subsequent steps such that v_i is clustered and $e_j \in \pi^{mid}(v_i)$. The fact that the length of $\pi^{mid}(v_i)$ is bounded plays an important role in the subsequent steps.

The remaining set of replacement paths are handled in Stage S3. This stage consists of two substages as well. In Substage S3.1, a collection of replacement paths $\{P^*_{i,j}\}$ that satisfies some key properties is constructed. Then, in Substage S3.2, the algorithm employs a modified path-buying procedure, first developed by Baswana et al. [4] and recently revisited by Cygan et al. [10]. This modified path-buying procedure heavily exploits the key properties of the paths $P^*_{i,j}$ constructed in Substage S3.1.

We now provide a detailed description of the algorithm and establish Theorem 4.1.

4.1 S1.1: Clustering

The following fact is taken from [10].

Fact 4.3 ([10]) *There is a poly-time algorithm* $\mathsf{Cluster}(G, \gamma)$ *that given a parameter* $\gamma \in [0,1]$ *and a graph* $G = (V, E)$ *constructs a collection* $\mathcal{C}$ *of at most* $n^{1-\gamma}$ *vertex-disjoint clusters, each of size* n^{γ}*, and a subgraph* G_C *of* G *with* $O(n^{1+\gamma})$ *edges, such that*

(1) for any missing edge $(u,v) \in E(G) \setminus E(G_C)$*,* u *and* v *belong to two different clusters, and*

(2) the diameter of each cluster (i.e., the maximum distance in G_C *between any two vertices of the cluster) is at most* 2*.*

Note that the clustering does not necessarily form a partition of V, i.e., the set $\widetilde{V}_C = V \setminus \bigcup_{C_i \in \mathcal{C}} C_i$ may be nonempty. However, in that case, all edges connecting vertices

of $\widetilde{V}_C$ among themselves or to vertices of one of the clusters belong to G_C. Also, the clusters are not guaranteed to be connected.

Invoke the algorithm $\mathsf{Cluster}$ on $G \setminus T_0$ and $\gamma = 1/3$, getting the subgraph $G_C = \mathsf{Cluster}(G \setminus T_0, \gamma)$ and let $\mathcal{C} = \{C_1, \ldots, C_\kappa\}$ be the clusters of G_C. Define

$$\mathcal{F}_1 = T_0 \cup G_C.$$

By Fact 4.3, $|E(\mathcal{F}_1)| = O(n^{4/3})$ and the number of clusters is $\kappa = O(n^{2/3})$. For a clustered vertex $v' \notin \widetilde{V}_C$, denote by $C(v') \in \mathcal{C}$ its cluster in G_C. For every cluster $C_\ell \in \mathcal{C}$, let $\mathtt{LCA}(C_\ell)$ be the least common ancestor of all C_ℓ vertices in T_0. Formally, $\mathtt{LCA}(C_\ell)$ is the vertex $\widehat{v}$ of maximal depth satisfying that $\pi(s, \widehat{v}) \subseteq \pi(s, v')$ for every $v' \in C_\ell$.

4.2 S1.2: Shortest path $\pi(s, v_i)$ segmentation

For every clustered vertex $v \in C_k$, divide its shortest path $\pi(s, v)$ into 3 segments

$$\pi(s, v) = \pi^{far}(v) \circ \pi^{mid}(v) \circ \pi^{near}(v)$$

in the following manner. Define $z_k \in \pi(s, \mathtt{LCA}(C_k))$ to be the upmost vertex on $\pi(s, \mathtt{LCA}(C_k))$ (closest to s) satisfying that $\mathrm{dist}(z_k, \mathtt{LCA}(C_k)) \leq \lceil n^{2/3} \rceil$. Then, let $\pi^{far}(v) = \pi(s, z_k)$, $\pi^{mid}(v) = \pi(z_k, \mathtt{LCA}(C_k))$ and $\pi^{near}(v) = \pi(\mathtt{LCA}(C_k), v)$. Note that for every two vertices $v, v' \in C_\ell$ in the same cluster it holds that $\pi^{mid}(v) = \pi^{mid}(v')$ and also $\pi^{far}(v) = \pi^{far}(v') \subseteq \pi(s, v)$. For an illustration of the segmentation of $\pi(s, v)$, see part S1.2 of Fig. 9; the vertex $\mathtt{LCA}(C_k)$ is used to draw the line between the near segment $\pi^{near}(v)$ and the rest of the $\pi(s, v)$ path. By the definition of $z_k \in \pi(s, \mathtt{LCA}(C_k))$, it holds that $|\pi^{mid}(v)| \leq \lceil n^{2/3} \rceil$.

Define $\pi(C_k) = \pi(s, \mathtt{LCA}(C_k))$ to be the maximal shortest path segment shared by the members of cluster C_k.

Observation 4.4 *(a)* $\pi(C_k) = \pi^{far}(v) \circ \pi^{mid}(v)$ *for every* $v \in C_k$.

(b) $\mathrm{dist}(v', v, G) \geq \lceil n^{2/3} \rceil$ *for every* $v' \in \pi^{far}(v)$.

Note that $\pi^{far}(v)$ is the same for every $v \in C_k$ and also $\pi^{mid}(v)$ is the same for every $v \in C_k$, therefore the common shortest-path section $\pi(C_k)$ can be divided into two segments $\pi^{far}(C_k)$ and $\pi^{mid}(C_k)$ such that

$$\pi^{mid}(C_k) \;=\; \pi^{mid}(v) \text{ for every } v \in C_k \text{ , and } \pi^{far}(C_k) = \pi^{far}(v) \text{ for every } v \in C_k \text{ .}$$

4.3 S2: Handling near and far edge faults

To protect against the edge failures occurring in the *near* and *far* shortest-path segments, two sparse collections of new edges are added to H, namely, E^{near} and E^{far}.

To guarantee that E^{near} and E^{far} are sufficiently sparse, the replacement paths are required to be *nice* as defined below. For a collection of replacement paths $\mathcal{Q} = \{Q^*_{i,j} \in SP(s, v_i, G \setminus \{e_j\})\}$, let $b_{i,j}$ be the first divergence point of $Q^*_{i,j} \in \mathcal{Q}$ from $\pi(s, v_i)$ (i.e., the first vertex on $Q^*_{i,j}$ such that the vertex v' appearing after it on $Q^*_{i,j}$ is not on $\pi(s, v_i)$). For every clustered vertex v_i, define

$$E^{near}_i \;=\; \{\texttt{LastE}(Q^*_{i,j}) \notin \mathcal{F}_1 \mid e_j \in \pi^{near}(v_i)\} \;\text{ and }\; E^{far}_i = \{\texttt{LastE}(Q^*_{i,j}) \notin \mathcal{F}_1 \mid e_j \in \pi^{far}(v_i)\}\ .$$

Definition 4.5 *A collection $\mathcal{Q} = \{Q^*_{i,j}\}$ of replacement paths is* nice *if it satisfies the following properties:*

(N1) $|E^{near}_i| \leq 5$,

(N2) $|E^{far}_i| \leq \lceil n^{1/3} \rceil$.

S2.1: The nice collection $\mathcal{Q}$ of replacement paths. Our starting point for constructing the nice collection $\mathcal{Q} = \{Q^*_{i,j} \in SP(s, v_i, G \setminus \{e_j\})\}$ is the graph $\mathcal{F}_1 = G_C \cup T_0$. Note that by Fact 4.3, the edges in $G \setminus \mathcal{F}_1$ are those incident to clustered vertices. A path $Q^*_{i,j} \in SP(s, v_i, G \setminus \{e_j\})$ is *cute* if its divergence point $b_{i,j}$ is unique (i.e., $Q^*_{i,j}[b_{i,j}, v_i]$ and $\pi(b_{i,j}, v_i)$ are edge disjoint). The following algorithm constructs a cute collection of replacement paths which are also shown to be nice.

Algorithm `Qcons` for $Q^*_{i,j}$ construction. For every vertex-edge pair (i, j) such that $v_i \in V(G)$ and $e_j \in \pi(s, v_i)$, define $Q^*_{i,j} \in SP(s, v_i, G \setminus \{e_j\})$ in the following manner. First, the replacement path is classified according to whether or not it *must* be missing ending (ending with an edge that is not in $\mathcal{F}_1$). To do that, the algorithm checks if there exists a replacement path which is not missing ending and sets $Q^*_{i,j}$ accordingly. This is done as follows. Let $E' = E(v_i, G \setminus E(\mathcal{F}_1))$ be the new edges incident to v_i in G.
Case (a): $\text{dist}(s, v_i, G \setminus (E' \cup \{e_j\})) = \text{dist}(s, v_i, G \setminus \{e_j\})$. In this case, there is a replacement path which is not missing-ending, and the algorithm takes one such path $Q^*_{i,j} \in SP(s, v_i, G \setminus (E' \cup \{e_j\})))$.
Case (b): $\text{dist}(s, v_i, G \setminus (E' \cup \{e_j\})) > \text{dist}(s, v_i, G \setminus \{e_j\})$. In this case, the chosen replacement path must be missing-ending. The algorithm attempts to select a *cute* replacement whose divergence point is highest. Let $\mathcal{U} = \{P \in SP(s, v_i, G \setminus \{e_j\}) \mid P \text{ is cute}\}$ be the collection of $s - v_i$ replacement paths which are cute. In the analysis section, we show that $\mathcal{U}$ is nonempty. For every cute path $P_\ell \in \mathcal{U}$ with unique divergence point b_ℓ, let the cost of the path be the depth (distance from s in T_0) of b_ℓ, i.e., $\texttt{Cost}(P) = \texttt{depth}(b_\ell)$. Let

$Q^*_{i,j} \in \mathcal{U}$, be the cute path of minimum cost. I.e., $\texttt{Cost}(Q^*_{i,j}) = \min\{\texttt{Cost}(P), P \in \mathcal{U}\}$.

Analysis of the $Q^*_{i,j}$ paths. In this section we prove the following.

Lemma 4.6 *$\mathcal{Q} = \{Q^*_{i,j}\}$ is nice.*

We begin by proving correctness.

Claim 4.7 *For every i, j,*

(a) $Q^*_{i,j} \in SP(s, v_i, G \setminus \{e_j\})$,

(b) if $\texttt{LastE}(Q^*_{i,j}) \notin T_0$ then $\text{dist}(s, v_i, G \setminus (E' \cup \{e_j\})) > \text{dist}(s, v_i, G \setminus \{e_j\})$ for $E' = E(v_i, G) \setminus T_0$.

Proof: If the path $Q^*_{i,j}$ is chosen by Algorithm `Qcons` according to case (a), then it is not a missing-ending and the lemma follows trivially. Note that Part (b) is also immediate by the construction. It remains to prove Claim (a) for paths $Q^*_{i,j}$ chosen according to Case (b), i.e., missing-ending $Q^*_{i,j}$ paths. In fact, since Algorithm `Qcons` chooses the cute $s - v_i$ replacement path whose divergence point is closest to s, it is sufficient to show that the set of cute paths $\mathcal{U} \subseteq SP(s, v_i, G \setminus \{e_j\})$ is nonempty. To do so, we exhibit at least one such path P' in this set. Let $P_1 \in SP(s, v_i, G \setminus \{e_j\})$ be an arbitrary replacement path. P_1 is converted into a *cute* path $P' \in \mathcal{U}$. Let $w \in (\pi(s, v_i) \cap P_1) \setminus \{v_i\}$ be the last mutual point of P_1 and $\pi(s, v_i)$ which is not v_i. Define $P' = \pi(s, w) \circ P_1[w, v_i]$. Clearly, $|P'| = |P_1|$, so it remains to show that $P' \subseteq (G \setminus \{e_j\})$ and in particular it is sufficient to show that $e_j \notin \pi(s, w)$. Let $\widehat{Q}_1 = \pi(s, w)$ and $\widehat{Q}_2 = \pi(w, v_i)$. We now claim that $e_j \in \widehat{Q}_2$. Assume, towards contradiction that $e_j \notin \widehat{Q}_2$. This implies that the path $Q' = P_2[s, w] \circ \widehat{Q}_2$ satisfies $|Q'| = |P_2|$ and $Q' \subseteq G \setminus \{e_j\}$, hence $Q' \in SP(s, v_i, G \setminus \{e_j\})$ and it is *not* missing-ending, in contradiction to the fact that $Q^*_{i,j}$ was chosen according to Case (b). Hence, $e_j \in \widehat{Q}_3$ and $P' \subseteq G \setminus \{e_j\}$ is cute. ∎

We next provide several preliminary claims.

Claim 4.8 *If $Q^*_{i,k}$ and $Q^*_{i,k'}$ are the missing-ending replacement paths chosen by Algorithm `Qcons` and $|Q^*_{i,k}| < |Q^*_{i,k'}|$, then e_k is below $e_{k'}$ on $\pi(s, v_i)$.*

Proof: If e_k is above $e_{k'}$, then $Q^*_{i,k} \in G \setminus \{e'_k\}$ so it can serve as a replacement path for v_i and $e_{k'}$ as well, in contradiction to Cl. 4.7, by which $Q^*_{i,k'}$ is an $s - v_i$ shortest path in $G \setminus \{e'_k\}$. ∎

Claim 4.9 *If $Q^*_{i,k}$ and $Q^*_{i,k'}$ are the missing-ending replacement paths chosen by Algorithm `Qcons` and $|Q^*_{i,k}| = |Q^*_{i,k'}|$, then $Q^*_{i,k} = Q^*_{i,k'}$.*

Proof: Without loss of generality, let e_k be above $e_{k'}$ on $\pi(s,v_i)$ and let b_k (resp., $b_{k'}$) be the unique divergence point of $\pi(s,v_i)$ and $Q^*_{i,k}$ (resp., $Q^*_{i,k'}$). By construction, there exists such unique divergence point. Let b_ℓ for $\ell \in \{k,k'\}$ be the divergence point that is closer to s. Since b_k must appear on $\pi(s,v_i)$ above e_k, it holds that b_ℓ appears above e_k as well and hence that $Q^*_{i,\ell} \in G \setminus \{e_k, e_{k'}\}$. In addition, since $|Q^*_{i,\ell}| = |Q^*_{i,k}| = |Q^*_{i,k'}|$ it holds that $Q^*_{i,\ell} \in SP(s,v_i,G\setminus\{e_k,e_{k'}\})$. By Algorithm `Qcons` definition, it holds that $b_{k'} = b_k$ and by the uniqueness of the shortest-path under W it holds that $Q^*_{i,k} = Q^*_{i,k'}$. ∎

We now establish the niceness of $\mathcal{Q}$. Let us first prove property (N1) of Def. 4.5.

Lemma 4.10 $|E^{near}_i| \leq 5$.

Proof: Assume towards contradiction that there exists some v_i such that $|E^{near}_i| \geq 6$. Let us consider 6 specific edges $e'_1,\ldots,e'_6 \in E^{near}_i$. Recall that v_i must be clustered and let $Q^*_{i,1},...,Q^*_{i,5},Q^*_{i,6}$ be replacement paths whose last edge is $e'_1,\ldots,e'_6$ respectively. It then holds that $e_j \in \pi^{near}(v_i) = \pi(w,v_i)$ where $w = \texttt{LCA}(C(v_i))$ is the least common ancestor of v_i's cluster $C(v_i)$. Since $\texttt{LastE}(Q^*_{i,j_1}) \neq \texttt{LastE}(Q^*_{i,j_2})$ for every $j_1,j_2 \in \{1,\ldots,6\}$, by Cl. 4.9, the paths $Q^*_{i,j}$ are of distinct lengths for every $j \in \{1,\ldots,6\}$. Without loss of generality, assume that $|Q^*_{i,1}| < ... < |Q^*_{i,6}|$. We therefore have that

$$|Q^*_{i,6}| > |Q^*_{i,1}| + 4\ . \tag{9}$$

In addition by Cl. 4.8, it holds that the edges e_j are sorted in decreasing distance from s, i.e., $\text{dist}(s,e_1) > \ldots > \text{dist}(s,e_6)$. Since $e_1,\ldots,e_6 \in \pi(w,v_i)$ appear strictly below w, it holds that there exists at least one vertex $v' \in C(v_i)$ such that $e_6 = (x,y) \notin \pi(s,v')$ and hence also $e'_1,\ldots,e'_5 \notin \pi(s,v')$. See Fig. 10 for an illustration. This holds since otherwise, if e_6 belongs to $\pi(s,v'')$ for every $v'' \in C(v_i)$ then the vertex y is a common ancestor of the cluster vertices $C(v_i)$ and it is deeper then w, in contradiction to the fact that w is the least common ancestor). Hence by the cluster diameter property of Fact 4.3, it holds that

$$\begin{aligned} |Q^*_{i,6}| &= \text{dist}(s,v_i,G\setminus\{e_6\}) \leq \text{dist}(s,v',G\setminus\{e_6\}) + 2 = |\pi(s,v')| + 2 \\ &= \text{dist}(s,v',G\setminus\{e_1\}) + 2 \leq \text{dist}(s,v_i,G\setminus\{e_1\}) + 4 = |Q^*_{i,1}| + 4\ , \end{aligned}$$

which is in contradiction with Eq. (9). The lemma follows. ∎

Finally, we turn to establish Property (N2) of Def. 4.5 and thus establish Lemma 4.6. For every path P, let $\texttt{Cost}_W(P) = \sum_{e\in P} W(e)$.

Claim 4.11 $|E^{far}_i| \leq \lceil n^{1/3} \rceil$.

Proof: We first claim that for every two missing-ending paths $Q^*_{i,j}$ and $Q^*_{i,j'}$ such that $\texttt{LastE}(Q^*_{i,j}) \neq \texttt{LastE}(Q^*_{i,j'}) \in E^{far}_i$, it holds that their detours $Q_1 = Q^*_{i,j}[b_{i,j},v_i]$ and

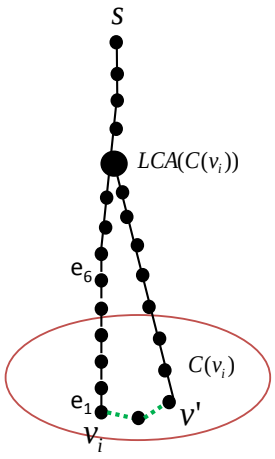

Figure 10: The set of E_i^{near} must contain at most 6 edges. Drawn are the edges $e_1, \ldots, e_6$ whose corresponding last edge of $Q^*_{i,1}, \ldots Q^*_{1,6}$ are in E_i^{near}. The edges of the cluster graph G_C are dotted. The vertices v_i and v' are in the same cluster hence of distance 2 in G_C.

$Q_2 = Q^*_{i,j'}[b_{i,j'}, v_i]$ are vertex disjoint except for their common endpoint v_i. Assume towards contradiction that there exists some mutual point $w \in Q_1 \cap Q_2 \setminus \{v_i\}$ in the intersection. Since Q_ℓ and $\pi(s, v_i)$ are edge disjoint for $\ell \in \{1, 2\}$, we have that there are two alternative $w - v_i$ paths in $G \setminus \{e_j, e_{j'}\}$, namely $Q_1[w, v_i]$ and $Q_2[w, v_i]$. By the optimality of $Q_1 \in SP(b_{i,j}, v_i, G \setminus \pi(s, v_i), W)$, we have that $\texttt{Cost}_W(Q_1[w, v_i]) < \texttt{Cost}_W(Q_2[w, v_i])$. Similarly, by the optimality of $Q_2 \in SP(b_{i,j'}, v_i, G \setminus \pi(s, v_i), W)$, we have that $\texttt{Cost}_W(Q_2[w, v_i]) < \texttt{Cost}_W(Q_1[w, v_i])$; contradiction. It follows that Q_1 and Q_2 are vertex disjoint. We next claim that it also holds that $|Q^*_{i,j}[b_{i,j}, v_i]| \geq \lceil n^{2/3} \rceil$. To see this, note that since $Q^*_{i,j}$ protects the fault of an edge $e_j \in \pi^{far}(v_i)$ and the unique divergence point $b_{i,j}$ must appear above e_j, it holds that $b_{i,j} \in \pi^{far}(v_i)$ as well. Hence, $|Q^*_{i,j}[b_{i,j}, v_i]| \geq \text{dist}(b_{i,j}, v_i, G) > \lceil n^{2/3} \rceil$ where the last inequality follows from Obs. 4.4(b).

Assume towards contradiction that $|E_i^{far}| > \lceil n^{1/3} \rceil$. Then there are $\lceil n^{1/3} \rceil$ vertex disjoint paths in G each of length at least $\lceil n^{2/3} \rceil$. Overall the number of vertices in those paths is greater than $|V(G)| = n$, contradiction. ∎

We conclude this section with the following immediate observation.

Observation 4.12 *For every clustered vertex $v_i \in V(G)$ and edge $e_j \in \pi(s, v_i)$, if $\texttt{LastE}(Q^*_{i,j}) \notin E_i^{near} \cup E_i^{far}$ then $e_j \in \pi^{mid}(v_i) = \pi^{mid}(C(v_i))$.*

S2.2: Creating $\mathcal{F}_2$. Having the nice collection of replacement paths $\mathcal{Q}$, let $IH = \{i \mid v_i \in V(G)\}$. The current spanner of this step is given by

$$\mathcal{F}_2 = \mathcal{F}_1 \cup \bigcup_{i \in IH} \left(E_i^{near} \cup E_i^{far} \right).$$

Note that by Definition 4.5, $|\mathcal{F}_2| = O(n^{4/3})$.

4.4 S3: Handling mid edge faults

At this point, we handled faults of edges $e_j \in \pi(s, v_i)$ except those occurring on the middle sections $\pi^{mid}(v_i)$. Unfortunately, those cannot be handled by adding the last edges of all missing-ending paths $Q^*_{i,j}$, as there may be too many such missing edges. Instead, we would like to "aggregate" the remaining problems, and handle them by adding only *some* of the missing edges relying on the properties of the clustering to provide approximately shortest replacement paths whenever the optimal path was not included. One property that could have helped this aggregation process is *prefix consistency*. The replacement paths $P^*_{i,j}$ and $P^*_{i',j}$ for $v_{i'} \in P^*_{i,j}$ are prefix-consistent if $P^*_{i',j} = P^*_{i,j}[s, v_{i'}]$. The advantage of this property is that by adding (missing) last edges of $P^*_{i',j}$, we also help the longer path $P^*_{i,j}$, see Fig. 11. Unfortunately, our construction of the path collection $\mathcal{Q}$ does not guarantee this property. This is the main motivation for the next step in the algorithm, where we replace the paths $\mathcal{Q}$ by a new path collection $P^*_{i,j}$ which is somewhat closer to achieving this property (albeit it falls short of doing that.)

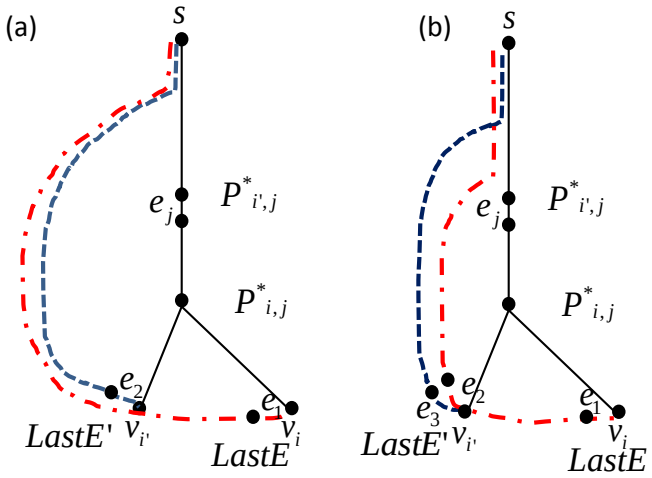

Figure 11: Illustration of the prefix consistency property. (a) The paths $P^*_{i,j}$ and $P^*_{i',j}$ are prefix consistent since $P^*_{i,j}[s, v_{i'}] = P^*_{i',j}$. (b) The paths $P^*_{i,j}$ and $P^*_{i',j}$ are not prefix consistent since $P^*_{i,j}[s, v_{i'}] \neq P^*_{i',j}$

We first describe the algorithm and then prove that these replacement paths satisfy some important properties which are crucial for the efficiency of the subsequent path-buying procedure.

S3.1: Algorithm `Pcons` for Constructing $P^*_{i,j}$. Consider the edges of T_0 in *nonincreasing* distance from s. For edge $e_j \in T_0$, let the set of vertices sensitive to the failure of e_j, $\mathcal{S}(e_j) = \{v_1, \ldots, v_k\}$, be sorted in *nondecreasing* distance from s in $G \setminus \{e_j\}$. I.e., $\text{dist}(s, v_1, G \setminus \{e_j\}) \leq \ldots \leq \text{dist}(s, v_k, G \setminus \{e_j\})$. For every e_j, the algorithm constructs the replacement path $P^*_{i,j}$ for vertices $v_i \in \mathcal{S}(e_j)$ according to this order. Initially set $P^{all}_i = \emptyset$ and $P^*_{i,j} = \emptyset$ for every $v_i \in V(G)$ and $e_j \in \pi(s, v_i)$. The algorithm constructs two intermediate replacement paths, $P^1_{i,j}$ and $P^2_{i,j}$. The final replacement path $P^*_{i,j}$ is obtained from $P^2_{i,j}$.

For vertex-edge pair (i, j), let v_ℓ be the neighbor of v_i on $Q^*_{i,j}$, i.e., $(v_\ell, v_i) = \texttt{LastE}(Q^*_{i,j})$ and do the following. Consider the following cases.
(R1.1) If $e_j \notin \pi(s, v_\ell)$, define $P^1_{i,j} = \pi(s, v_\ell) \circ (v_\ell, v_i)$. Note that by construction, $P^1_{i,j}$ and $Q^*_{i,j}$ share the same last edge.
(R1.2) If $e_j \in \pi(s, v_\ell)$ (i.e., $v_\ell \in \mathcal{S}(e_j)$) let $P^1_{i,j} = P^*_{\ell,j} \circ (v_\ell, v_i)$. (Note that $P^*_{\ell,j}$ was already constructed, since v_ℓ was handled before v_i in our ordering.) Also note that the path collection $P^1_{i,j}$ does achieve the desired property of prefix consistency. Unfortunately, this set no longer enjoys another desirable property that possessed by the original collection $\mathcal{Q}$, namely, the uniqueness of the divergence points $b_{i,j}$. This is why we next modify the paths further, defining the paths $P^2_{i,j}$ which restores the uniqueness property while possibly destroying prefix consistency again.
(R2.1) If $(v_\ell, v_i) \in \mathcal{F}_2$ is *not missing*, then let $P^*_{i,j} = P^2_{i,j} = P^1_{i,j}$ and proceed to the next vertex-edge pair.
(R2.2) Else, $\texttt{LastE}(P^1_{i,j}) = (v_\ell, v_i) \notin \mathcal{F}_2$. (Note that this means that $\texttt{LastE}(P^1_{i,j}) \notin E^{near} \cup E^{far}$, and hence by Obs. 4.12, in this case v_i is clustered and $e_j \in \pi^{mid}(C(v_i))$.)

Let b be the first divergence point of $P^1_{i,j}$ and $\pi(s, v_i)$. If $P^1_{i,j}[b, v_i]$ and $\pi(b, v_i)$ are not edge disjoint (b is not a unique divergence point) then let $w \in \big(\pi(b, v_i) \cap P^1_{i,j}[b, v_i]\big) \setminus \{v_i\}$ be the last mutual vertex in $\pi(b, v_i)$ and $P^1_{i,j}[b, v_i]$ which is not v_i. Our goal now is to make w the unique divergence point with $\pi(s, v_i)$.
(R3) Defining $P^2_{i,j}$: $P^2_{i,j} = \pi(s, w) \circ P^1_{i,j}[w, v_i]$, making $b_{i,j} = w$ the unique divergence point of $P^2_{i,j}$ and $\pi(s, v_i)$.
(R4) Defining $P^*_{i,j}$: (R4.1) If $b_{i,j} \in \pi^{mid}(v_i)$, then set $P^*_{i,j} = P^2_{i,j}$ as the final replacement path. (R4.2) Otherwise, if $b_{i,j} \notin \pi^{mid}(v_i)$, do the following. First, if $P^{all}_i = \emptyset$, set $P^{all}_i \leftarrow P^2_{i,j}$. Finally, let $P^*_{i,j} \leftarrow P^{all}_i$. This completes the description of the algorithm.

Let

$$\mathcal{F}_3 = \mathcal{F}_2 \cup \bigcup \{\texttt{LastE}(P^{all}_i) \;\mid\; v_i \in V(G)\}.$$

The subgraph $\mathcal{F}_3$ contains, in addition to the edges of $\mathcal{F}_2$, also the last edges of P_i^{all} for every $v_i \in V(G)$. Since each such vertex contributes at most one edge from P_i^{all} to $\mathcal{F}_3$, at most n edges are added, hence $|\mathcal{F}_3| = O(n^{4/3})$.

Analysis of $P^*_{i,j}$ paths. We begin by showing that the constructed replacement paths $P^*_{i,j} \in SP(s, v_i, G \setminus \{e_j\})$ are of optimal lengths.

Lemma 4.13 *For every $v_i \in V(G)$ and $e_j \in \pi(s, v_i)$, consider $P^*_{i,j}$ and let $V^+_{i,j}$ be the set of vertices appearing on $P^*_{i,j}$ after the first new edge $e = \texttt{FirstNewE}(P^*_{i,j}) \notin T_0$. Then*

(a) $P^*_{i,j} \in SP(s, v_i, G \setminus \{e_j\})$ and

(b) $V^+_{i,j} \subseteq \mathcal{S}(e_j)$.

Proof: We prove this by induction on the iteration in which $P^*_{i,j}$ was constructed by the algorithm. For the induction base, consider the first iteration, when $P^*_{1,1}$ is constructed. By the ordering, e_1 is the last edge of some source s to leaf v path, i.e., $e_1 = \texttt{LastE}(\pi(s, v))$, and v_1 is the first vertex in the ordered set $\mathcal{S}(e_1)$. Let $(v_\ell, v_1) = \texttt{LastE}(Q^*_{i,1})$. We first prove (a) and show that $P^*_{1,1} \in SP(s, v_1, G \setminus \{e_1\})$. By Cl. 4.7 we have that $Q^*_{1,1} \in SP(s, v_1, G \setminus \{e_1\})$ and therefore $\text{dist}(s, v_\ell, G \setminus \{e_\ell\}) < \text{dist}(s, v_1, G \setminus \{e_1\})$. As v_1 is the first vertex in the ordering of $\mathcal{S}(e_1)$ it implies that $v_\ell \notin \mathcal{S}(e_1)$. Since $\pi(s, v_\ell)$ has at most one divergence point $b_{i,j}$ with $\pi(s, v_i)$ (in particular, $b_{i,j} = \texttt{LCA}(v_i, v_\ell)$), it holds that $P^2_{1,1} = P^1_{1,1}$ and since this is that first iteration P_i^{all} was not previously defined. Therefore, $P^*_{1,1} = P^1_{1,1} = P^2_{1,1} = \pi(s, v_\ell) \circ (v_\ell, v_1)$ is such that $|P^*_{1,1}| = |Q^*_{1,1}|$ and $P^*_{1,1} \subseteq G \setminus \{e_1\}$. Claim (a) is established. Consider Claim (b). By the definition of $P^*_{1,1}$, it holds that $P^*_{1,1} \setminus \{(v_\ell, v_1)\} \subseteq T_0$, hence $V^+_{1,1} = \{v_1\}$, and since $v_1 \in \mathcal{S}(e_1)$, the induction base holds.

Assume Claims (a) and (b) of the lemma hold for every replacement path $P^*_{i',j'}$ constructed up to iteration $t-1$ and consider the path $P^*_{i,j}$ constructed at iteration t. Let $(v_\ell, v_i) = \texttt{LastE}(P^*_{i,j})$. We first establish the lemma for $P^1_{i,j}$, then for $P^2_{i,j}$, and finally consider the case where $P^*_{i,j} \neq P^2_{i,j}$, i.e., where $P^*_{i,j} = P_i^{all}$. If $v_\ell \notin \mathcal{S}(e_j)$, then Algorithm `Pcons` again yields $P^*_{i,j} = P^2_{i,j} = P^1_{i,j} = \pi(s, v_\ell) \circ (v_\ell, v_i)$, and similarly to the induction base, by Cl. 4.7, we have that $P^*_{i,j} \in SP(s, v_i, G \setminus \{e_j\})$, yielding Claim (a), and since $P^*_{i,j} \setminus \{\texttt{LastE}(P^*_{i,j})\} \subseteq T_0$, Claim (b) holds as well. For the rest of the proof, it remains to consider the case where $v_\ell \in \mathcal{S}(e_j)$. By Cl. 4.7, $Q^*_{i,j} \in SP(s, v_i, G \setminus \{e_j\})$, hence $\text{dist}(s, v_\ell, G \setminus \{e_j\}) < \text{dist}(s, v_i, G \setminus \{e_j\})$ and by the ordering of $\mathcal{S}(e_j)$, the pair (ℓ, j) was considered at iteration $t' < t$ and the induction assumption for $P^*_{\ell,j}$ can be applied. We then have that $P^1_{i,j} = P^*_{\ell,j} \circ \texttt{LastE}(Q^*_{i,j}) \in SP(s, v_i, G \setminus \{e_j\})$, which satisfies Claim (a).

We now show that Claim (b) holds for the path $P^1_{i,j}$. Note that since $v_\ell \in \mathcal{S}(e_j)$, the path $P^*_{\ell,j}$ must contain a new edge, so $\texttt{FirstNewE}(P^1_{i,j}) = \texttt{FirstNewE}(P^*_{\ell,j})$. Then, by the

induction assumption for $P^*_{\ell,j}$ and the fact that $v_i \in \mathcal{S}(e_j)$, we have that $V(P^1_{i,j}[y, v_i]) \subseteq \mathcal{S}(e_j)$ where $\texttt{FirstNewE}(P^*_{\ell,j}) = (x, y)$. Hence, so far, the lemma holds for $P^1_{i,j}$. It remains to consider the cases where $P^*_{i,j} \neq P^1_{i,j}$, and hence $\texttt{LastE}(P^1_{i,j}) \notin \mathcal{F}_2$.

We next show that the lemma holds for $P^2_{i,j}$. Since $\texttt{LastE}(P^1_{i,j}) = \texttt{LastE}(Q^*_{i,j}) = (v_\ell, v_i) \notin T_0$, Cl. 4.7(b) implies that *any* $s - v_i$ replacement path in $G \setminus \{e_j\}$ must be *missing-ending*. Let b be the first divergence point of $P^1_{i,j}$ and $\pi(s, v_i)$. If $P^2_{i,j} \neq P^1_{i,j}$ then necessarily there exists another mutual point $w \in \big(\pi(b, v_i) \cap P^1_{i,j}[b, v_i]\big) \setminus \{b, v_i\}$ such that $P^2_{i,j} = \pi(s, w) \circ P^1_{i,j}[w, v_i]$. Consider Claim (a). Since $P^1_{i,j} \subseteq G \setminus \{e_j\}$ and $|P^2_{i,j}| \leq |P^1_{i,j}|$, to establish Claim (a), it is sufficient to show that $e_j \notin \pi(s, w)$.

Observation 4.14 $e_j \in \pi(w, v_i)$.

Proof: Assume, towards contradiction, that $e_j \notin \pi(w, v_i)$. This implies that the path $P' = P^1_{i,j}[s, w] \circ \pi(w, v_i)$ satisfies $|P'| = |P^1_{i,j}|$ and $P' \subseteq G \setminus \{e_j\}$, hence $P' \in SP(s, v_i, G \setminus \{e_j\})$ and it is *not* missing-ending as $\texttt{LastE}(P') \in E(T_0)$, contradicting Cl. 4.7(b). ∎

Hence, $e_j \notin \pi(w, v_i)$, concluding that $P^2_{i,j} \in SP(s, v_i, G \setminus \{e_j\})$ as required, so Claim (a) of the lemma holds for $P^2_{i,j}$. Consider Claim (b). The prefix $P^2_{i,j}[s, w]$ of $P^2_{i,j}$ is entirely in T_0 (since it is equal to $\pi(s, w)$), so $\texttt{FirstNewE}(P^2_{i,j})$ must occur in $P^2_{i,j}[w, v_i]$. Since $P^2_{i,j}[w, v_i] = P^1_{i,j}[w, v_i]$, the validity of Claim (b) for $P^1_{i,j}$ implies that Claim (b) holds for $P^2_{i,j}$ as well. Hence if $P^*_{i,j} = P^2_{i,j}$ then we are done.

It remains to consider the case where $P^*_{i,j} \neq P^2_{i,j}$ or in other words $P^*_{i,j} = P^{all}_i$. Let $t' < t$ be the iteration in which P^{all}_i was defined and let (i, j') be the pair considered at iteration t', hence $P^{all}_i = P^*_{i,j'} = P^2_{i,j'}$. By the induction assumption for $t' < t$, Claims (a) and (b) hold for $P^*_{i,j'}$. Note that in this case, it holds that $b_{i,j}$ (resp., $b_{i,j'}$), the unique divergence point of $P^2_{i,j}$ (resp., $P^2_{i,j'}$) and $\pi(s, v_i)$, is not in $\pi^{mid}(C(v_i))$.

Observation 4.15 $b_{i,j}, b_{i,j'} \in \pi^{far}(v_i)$.

Proof: By the structure of the algorithm (step (R2.2)), it holds that $\texttt{LastE}(Q^*_{i,j}), \texttt{LastE}(Q^*_{i,j'}) \notin \mathcal{F}_2$ and hence by Obs. 4.12 it holds that

$$e_j, e_{j'} \in \pi^{mid}(v_i) \ . \tag{10}$$

By the correctness established for the paths $P^2_{i,j'}$ and $P^2_{i,j}$ (Claim (a)), we have that $P^2_{i,j} \in SP(s, v_i, G \setminus \{e_j\})$ and $P^2_{i,j'} \in SP(s, v_i, G \setminus \{e_{j'}\})$. Hence, the unique divergence point $b_{i,j}$ (resp.,$b_{i,j'}$) of $P^2_{i,j}$ (resp., $P^2_{i,j'}$) and $\pi(s, v_i)$, appears on $\pi(s, v_i)$ above the failed edge e_j (resp., $e_{j'}$). Combining Eq. (10) with the fact that $b_{i,j'}, b_{i,j} \notin \pi^{mid}(v_i)$, it follows that $b_{i,j'}, b_{i,j}$ appear above the vertices of $\pi^{mid}(v_i)$ hence they appear on $\pi^{far}(v_i)$. ∎

By Claim (a) of the inductive assumption for $P^*_{i,j'} = P^{all}_i$, it holds that $P^{all}_i \in SP(s, v_i, G \setminus \{e_{j'}\})$. By the proof of Claim (a) of the lemma for $P^2_{i,j}$, it holds that

$P^2_{i,j} \in SP(s, v_i, G \setminus \{e_j\})$. Combining with Obs. 4.15, it holds that there are two replacement paths $P^2_{i,j}$ and P^{all}_i in $G \setminus \{e_j, e_{j'}\}$. By the optimality of P^{all}_i, it holds that $|P^{all}_i| \leq |P^2_{i,j}|$, and by the optimality of $P^2_{i,j}$, it holds that $|P^2_{i,j}| = |P^{all}_i|$. Claim (a) of the lemma holds. Consider Claim (b). By Claim (b) of the induction assumption for $P^*_{i,j'}$ it holds that $V^+_{i,j'} \subseteq \mathcal{S}(e_{j'})$. Since $e_{j'}$ is below e_j (as the edges of T_0 are considered in non increasing distance from s and $e_j, e_{j'} \in \pi(s, v_i)$) it holds that $V^+_{i,j} = V^+_{i,j'} \subseteq \mathcal{S}(e_{j'}) \subseteq \mathcal{S}(e_j)$. Claim (b) follows. The lemma holds. ∎

For a missing ending replacement path $P^*_{i,j}$, let $V^i_j = \{v^i_1, \ldots, v^i_r = v_i\}$ be the end-vertices of missing edges on $P^*_{i,j}$, i.e., $\texttt{LastE}(P^*_{i,j}[s, v^i_k]) \notin \mathcal{F}_3$ for every $v^i_k \in V^i_j$. Hereafter, let $b_{i,j}$ be the first divergence point of $P^*_{i,j}$ and $\pi(s, v_i)$.

Lemma 4.16 *The following properties hold for every $v^i_k \in V^i_j$:*

(a) $e_j \in \pi^{mid}(C(v^i_k))$ (hence $e_j \in \pi(s, v')$ for every $v' \in C(v^i_k)$),

(b) $b_{i,j}$ is the unique *divergence point of $\pi(s, v^i_k)$ and $P^*_{i,j}[s, v^i_k]$, hence $P^*_{i,j}[b_{i,j}, v^i_k]$ and $\pi(s, v^i_k)$ are edge disjoint, and*

(c) $b_{i,j} \in \pi^{mid}(C(v^i_k))$.

Proof: First note that since the vertices of V^i_j have missing edges in $\mathcal{F}_3$, it follows by Fact 4.3(2) that each of them is clustered. We prove the lemma by induction on the iteration t in which $P^*_{i,j}$ was constructed. For the induction base, consider $t = 1$ and let v_1 be the first vertex in $\mathcal{S}(e_1)$ where e_1 be the last edge of some s to leaf v path and consider $P^*_{1,1}$. Let $(v_\ell, v_1) = \texttt{LastE}(P^*_{1,1})$. By the induction base of Lemma 4.13, it holds that

$$P^*_{1,1} = \pi(s, v_\ell) \circ (v_\ell, v_1) \ . \tag{11}$$

Hence the only missing edge on $P^*_{1,1}$ is at most $\texttt{LastE}(P^*_{1,1})$. I.e., $V^1_1 = \emptyset$ or $V^1_1 = \{v_1\}$. If $V^1_1 = \emptyset$, the claim holds vacuously. So consider the case where $V^1_1 = \{v_1\}$ (i.e., $\texttt{LastE}(P^*_{1,1})$ is missing). Note that $\texttt{LastE}(P^*_{1,1}) = \texttt{LastE}(Q^*_{1,1})$. Hence, by Obs. 4.12, $e_1 \in \pi^{mid}(C(v_1))$, so Claim (a) holds.

By Eq. (11), $b_{1,1} = \texttt{LCA}(v_\ell, v_1)$ is the unique divergence point of $P^*_{1,1}$ and $\pi(s, v_1)$. Thus Claim (b) holds.

Consider Claim (c). Recall that in this case $V^1_1 = \{v_1\}$ i.e., $\texttt{LastE}(Q^*_{1,1}) \notin \mathcal{F}_3$. If $b_{1,1} \notin \pi^{mid}(C(v_1))$, then by step (S4.2) of Algorithm `Pcons`, $P^*_{1,1} = P^{all}_1$. Since $\texttt{LastE}(P^{all}_1) \in \mathcal{F}_3$, we end with contradiction. The induction base holds.

Assume the claims hold for all replacement paths constructed up to iteration $t-1$ and let (i, j) be the pair considered at iteration t. We first prove the lemma for the case where $P^*_{i,j} = P^2_{i,j}$ and then consider the case where $P^*_{i,j} \neq P^2_{i,j}$.

Let $(v_\ell, v_i) = \texttt{LastE}(Q^*_{i,j})$. If $v_\ell \notin \mathcal{S}(e_j)$, then $P^*_{i,j} = \pi(s, v_\ell) \circ (v_\ell, v_i)$. The correctness follows as in the induction base. Thus, it remains to consider the complementary case where $v_\ell \in \mathcal{S}(e_j)$ and by construction, $P^1_{i,j} = P^*_{\ell,j} \circ (v_\ell, v_i)$. By Cl. 4.7, $\text{dist}(s, v_\ell, G \setminus \{e_j\}) < \text{dist}(s, v_i, G \setminus \{e_j\})$ and hence the induction assumption for $P^*_{\ell,j}$ can be applied.

Consider the vertices V^ℓ_j with missing edges on $P^*_{\ell,j} = P^1_{i,j} \setminus \{\texttt{LastE}(P^1_{i,j})\}$. Claim (a) holds for V^ℓ_j by the induction assumption for $P^*_{\ell,j}$. If $V^i_j = V^\ell_j$ we are done, else it holds that $V^i_j = V^\ell_j \cup \{v_i\}$. Since in this case $\texttt{LastE}(P^1_{i,j}) = \texttt{LastE}(Q^*_{i,j}) \notin \mathcal{F}_3$, it holds by Obs. 4.12 that $e_j \in \pi^{mid}(C(v_i))$. Hence, Claim (a) is established for V^i_j.

Consider Claim (b). By Claim (b) of the induction assumption, the first divergence point of $P^*_{\ell,j}[s, v_k]$ and $\pi(s, v_k)$, namely, $b_{\ell,j}$ is unique and common for every $v_k \in V^\ell_j$. Let b be the first divergence point of $P^1_{i,j}$ and $\pi(s, v_i)$. We first claim the following.

Claim 4.17 $b = b_{\ell,j}$.

Proof: By Claim (a) of the induction assumption for $P^*_{\ell,j}$ it holds that $e_j = (x, y) \in \pi(s, v_k)$ for every $v_k \in V^\ell_j$. By definition, it also holds that $e_j \in \pi(s, v_i)$. Hence the first divergence point b satisfies $b \in P^*_{\ell,j}$. Let b' be the vertex that appears after b on $P^*_{i,j}$ (hence also on $P^*_{\ell,j}$). As b is a divergence point, it must hold that $b' \notin \pi(s, x)$ and therefore $b' \notin \pi(s, v_k)$ for every $v_k \in V^i_j$ and in particular $b' \notin \pi(s, v_i)$. Hence $b_{\ell,j} = b$ as required. The claim follows. ∎

Let $w \in \big(\pi(b, v_i) \cap P^1_{i,j}[b, v_i]\big) \setminus \{v_i\}$ be the last divergence point of $\pi(s, v_i)$ and $P^1_{i,j}$. If $w = b$, then $b = b_{i,j}$ is the unique divergence point for every $v_k \in V^i_j$ and Claim (b) holds. It remains to consider the case where $b_{i,j} = w \neq b$, and hence the algorithm takes $P^2_{i,j} = \pi(s, w) \circ P^1_{i,j}[w, v_i]$. In this case, we show that $\pi[s, v_k]$ is edge disjoint with $P^*_{i,j}[w, v_k]$ for every $v_k \in V^i_j \setminus \{v_i\} = V^\ell_j$. I.e., since w is by the definition the unique divergence point for v_i we now want to show that we did not "ruin" this property for the "surviving" vertices with missing edges on $P^*_{\ell,j} \cap P^*_{i,j} = P^*_{\ell,j}[w, v_\ell]$. Note that since $P^*_{i,j}[s, b] = \pi(s, b)$, it holds that $\texttt{depth}(w) > \texttt{depth}(b)$. By Claim (b) of the induction assumption for $P^*_{\ell,j}$, it holds that for every $v_k \in V^\ell_j$, the paths $\pi(s, v_k)$ and $P^*_{\ell,j}[b, v_k]$ are edge disjoint (where $b = b_{\ell,j}$). In addition, note that by the correctness of $P^2_{i,j}$ in Lemma 4.13(a), it holds that w, the unique divergence point of $P^2_{i,j}$ and $\pi(s, v_i)$ appears on $\pi(s, v_i)$ above the failed edge e_j. By the induction assumption for Claim (a) on $P^*_{\ell,j}$, it holds that $e_j \in \pi(s, v_k)$ for every $v_k \in V^\ell_j$. Combining the last two observations, it follows that $w \in \pi(s, v_k)$ for every $v_k \in V^\ell_j$. We therefore have that $\pi(w, v_k) \subseteq \pi(s, v_k)$ and $P^*_{i,j}[w, v_k] = P^*_{\ell,j}[w, v_k] \subseteq P^*_{\ell,j}[b, v_k]$ are edge disjoint. Finally, since $P^*_{i,j}[s, w] = \pi(s, w)$, it holds that $P^*_{i,j}[w, v_k]$ and $\pi(s, v_k)$ are edge disjoint for every $v_k \in V^\ell_j$ as required. Claim (b) is established.

Consider Claim (c). We first show that the claim holds for every $v_k \in V^\ell_j \cap P^2_{i,j}$. By induction assumption for $P^*_{\ell,j}$, we have that $b_{\ell,j} \in \pi^{mid}(C(v_k))$. By the proof of Claim

(b), it holds that $\texttt{depth}(b_{\ell,j}) \leq \texttt{depth}(b_{i,j})$. In addition, by Claim (a) of the lemma, $e_j = (x', y') \in \pi^{mid}(C(v_k))$. Hence the entire π segment from b to x' satisfies $\pi(b_{\ell,j}, x') \subseteq \pi^{mid}(C(v_k))$. Since $b_{i,j}$ is the unique divergence point (by Claim (b)) it appears above e_j but not above $b_{\ell,j}$, hence $b_{i,j} \in \pi(b_{\ell,j}, x') \subseteq \pi^{mid}(C(v_k))$. We now prove Claim (c) for the case where $V_j^i \setminus V_j^\ell = \{v_i\}$. Assume towards contradiction that $b_{i,j} \notin \pi^{mid}(C(v_i))$. Then, by construction in this case $P^*_{i,j} = P_i^{all}$. Since $\texttt{LastE}(P_i^{all}) \in \mathcal{F}_3$, we end with contradiction to the fact that $v_i \in V_j^i$. Hence, the lemma follows for every (i,j) such that $P^*_{i,j} = P^2_{i,j}$.

Finally, we consider the complementary case where $P^*_{i,j} = P_i^{all} \neq P^2_{i,j}$. Let (i,j') be the vertex edge pair considered when P_i^{all} was first defined, thus $P_i^{all} = P^2_{i,j'} = P^*_{i,j'}$. Since $P^*_{i,j'}$ was defined before $P^*_{i,j}$, the induction assumption can be applied. Claims (b) and (c) for $P^*_{i,j}$ follow immediately by the induction assumption of Claims (b) and (c) for P_i^{all}. To see Claim (a), note that by the ordering of the edges, $e_{j'} = (x', y')$ must be below e_j on $\pi(s, v_i)$, in addition, by Lemma 4.13(a) it holds that $P_i^{all} \in SP(s, v_i, G \setminus \{e_j\})$, hence $b_{i,j'}$ appears above e_j. By part (c), $b_{i,j'} \in \pi^{mid}(C(v_i))$, hence $e_j \in \pi(b_{i,j'}, x') \subseteq \pi^{mid}(C(v_i))$, as required. The lemma follows. ∎

S3.2: Path-Buying procedure. With the collection of replacement path $P^*_{i,j}$ at hand, we are now ready to present the last step of the algorithm, a modified path-buying procedure, where the replacement path is added entirely to the spanner if it satisfies a particular cost to value balance.

Generally speaking, the high level approach of the path-buying technique is as follows. Recall that in the preliminary clustering sub-stage S1, the graph was condensed into clusters. There is a collection of $s_i - t_i$ paths, whose distance in the final H is required to approximate the distance in G by an additive factor. These paths are examined sequentially, where at step τ, a particular candidate path P_τ is considered to be added to the current spanner H_τ, resulting in $H_{\tau+1}$. The decision is made by assigning each candidate path P_τ a *cost* $\texttt{Cost}(P_\tau)$, corresponding to the number of path edges not already contained in the spanner H_τ, and a *value* $\texttt{Val}(P_\tau)$, measuring how much adding the path would help to satisfy the considered set of constraints on the pairwise distances. The candidate path P_τ is added to H_τ if its *value to cost ratio* is sufficiently large. Informally, if a path P_τ is added, then it implies that each of at least some fraction of its new edges $P_\tau \setminus H_\tau$ *contributes* to improving the inter-cluster distances in the current spanner H_τ. In the context of Stage S3.4 of our algorithm for FT-ABFS structures, we are given a collection of replacement paths $P^*_{i,j}$ where $e_j \in \pi^{mid}(v_i)$ and some preliminary sparse subgraph $\widehat{E}$ consisting of (at least) the edges of the BFS tree T_0, the edges of clustering graph G_C and the set $\bigcup_i E_i^{far} \cup \bigcup_i E_i^{near}$, containing the last edges of replacement paths $P^*_{i',j'}$ protecting against the failure of $e_{j'} \in \pi^{near}(v_{i'}) \cup \pi^{far}(v_{i'})$. By the preliminary explanation above (see Obs. 4.2) it is sufficient to consider only replacement

paths $P^*_{i,j}$ whose last edge is missing in $\widehat{E}$. These paths have a special structure. In particular, there is a unique divergence point $b_{i,j}$ where $P^*_{i,j}$ diverges from $\pi(s, v_i)$ and does not meet it again (i.e., $P^*_{i,j}[b_{i,j}, v_i]$ and $\pi(b_{i,j}, v_i)$ are edge disjoint). Since the common prefix $P^*_{i,j}[s, b_{i,j}] = \pi(s, b_{i,j})$ is contained in T_0, the buying procedure restricts attention only to the "detour" segment $P^-_{i,j} = P^*_{i,j}[b_{i,j}, v_i]$. The properties of the partial spanner $\mathcal{F}_3$ constructed so far guarantee that this detour $P^-_{i,j}$ is restricted to G, i.e., $P^-_{i,j} \subseteq G$, hence the size of the resulting construct would be bounded as a function of $n = |V(G)|$ as desired.

To gain some intuition regarding our modified path-buying technique, we review some of its principles and draw some differences between our setting and that of [4] and [10]. The analysis of the path-buying technique has two main ingredients. The first is the correctness ingredient (A1), where it is required to show that if an $s_i - t_i$ path P^* was *not* added to the current spanner H_t at time t, then there exists an alternative $s_i - t_i$ path P' in H_t satisfying that $|P'| \leq |P^*| + \beta$ for some constant integer $\beta > 0$. The second is the size ingredient (A2), where it is required to show that the number of edges added due to the paths that were bought by the procedure is bounded.

In the 6-additive construction of [4], the correctness ingredient (A1) was based upon the fact that if a path P^* was not added, then there exists a vertex $v_q \in P^*$ such that the pairwise $C(s_i) - C(v_q)$ and $C(t_i) - C(v_q)$ distances between clusters in H_t is smaller than that in P^*, namely, than $\mathrm{dist}(s_i, v_q, P^*)$ and $\mathrm{dist}(t_i, v_q, P^*)$ respectively. Similarly, in the subsetwise construction of [10], (A1) was established by noting that if a path P^* was not added, then there exists a vertex $v_q \in P^*$ such that the vertex to cluster $s_i - C(v_q)$ distance as well as the $t_i - C(v_q)$ distance in H_t are smaller than $\mathrm{dist}(s_i, v_q, P^*)$ and $\mathrm{dist}(t_i, v_q, P^*)$ respectively. In our setting, this argument becomes more delicate due to the possibility of failures which might render the existing bypasses already available in H_t useless. Hence, when considering a detour $P^-_{i,j}$ of a replacement path $P^*_{i,j}$ that was not added to the current spanner H_t, it is required to show that the inter-cluster bypasses in H_t do not contain the failed edge e_j and hence can safely be used in the surviving structure $H_t \setminus \{e_j\}$.

We now consider the second ingredient of the analysis (A2). Let $\mathcal{B}$ be the set of paths added to the spanner by the path-buying procedure. In both [4] and [10], (A2) is established based on the fact that if $P^* \in \mathcal{B}$, then its value satisfies $\mathtt{Val}(P^*) \geq c \cdot \mathtt{Cost}(P^*)$ for some constant $c \geq 1$. This implies that each of at least a constant fraction of the newly added edges on P^* contributes by decreasing some specific inter-cluster distances in the given spanner. Specifically, the value of P^* is the number of pairs (x, C) where C is a cluster and there exists a vertex $v \in C \cap P^*$ such that by adding P^* to the current spanner H_t, $\mathrm{dist}(x, C, P^*)$ is improved compared to that in H_t. In the setting of [4], x is a cluster (i.e., $x \in \{C(s_i), C(t_i)\}$) and in the setting of [10], x is a vertex (i.e., $x \in \{s_i, t_i\}$). For every $P^* \in \mathcal{B}$, its value (total number of (x, C)

pairs) is proportional to its cost $\texttt{Cost}(P^*)$. Therefore, to bound the number of edges, it is sufficient to bound the number of (x, C) pairs. This involves two steps: (A2.1) showing that the contribution due to a fixed pair (x, C) is bounded (or in other words, that a given pair (x, C) can contribute only a bounded number of times to the value of the paths in $\mathcal{B}$) and (A2.2) showing that the number of distinct (x, C) pairs is bounded. The combination of (A2.1) and (A2.2) bounds the size of the spanner. We now consider (A2.1) and (A2.2) separately. In both [4] and [10], (A2.1) is established by the cluster diameter property of Fact 4.3, which implies in this context, that every given pair (x, C) can contribute at most a constant number of times to the values of the paths in $\mathcal{B}$. Turning to (A2.2), in [4], the total number of distinct pairs corresponds to $|\mathcal{C}| \times |\mathcal{C}|$ since $x \in \mathcal{C}$. In comparison, in [10], since $x \in S$, there are a total of $O(|\mathcal{C}| \times |S|)$ pairs. Overall, the spanner size is bounded since the number of clusters $\mathcal{C}$, as well as the cardinality of S in the subsetwise variant of [10], are bounded.

We now contrast this with the situation in our setting of `FT-ABFS` structures. Part (A2.1) is no longer straightforward, since every replacement path $P^*_{i,j}$ added to the current spanner exists in a *different* graph $G \setminus \{e_j\}$, and therefore, in contrast to [4] and [10], the bounded diameter of the clusters is not sufficient, by itself, to establish (A2.1). Considering (A2.2), the approach of [4] can be adopted to bound to number of distinct pairs (x, C), by letting $x \in \{C(b_{i,j}), C(v_i)\}$, but this would result in an additive stretch of 6. To improve the additive stretch to 4, it is necessary to employ some *intermediate* compromise. We first impose a *direction* on the candidate paths, hence breaking the symmetry between the path endpoints. The direction is imposed by the source s. In particular, the two endpoints $b_{i,j}$ and v_i of each detour $P^-_{i,j}$ are treated in an asymmetric manner. The endpoint v_i is treated as a *cluster* in a similar manner to that of [4], that is, the value of $P^-_{i,j}$ counts the improvement of the inter-cluster distances between $C(v_i)$ and some $C(v_k)$ for $v_k \in P^-_{i,j}$. In contrast, the $b_{i,j}$ endpoint is treated as a *vertex*, as in [10]. Overall, the pairs (x, C) that contribute to the value of $P^-_{i,j}$ are of two types, where $x \in \{b_{i,j}, C(v_i)\}$. In the analysis section, it is shown that the number of contributions of each fixed pair (x, C) is bounded, and moreover, that for every cluster $C \in \mathcal{C}$, the number of distinct pairs in which it appears (of both types) is bounded by $O(n^{4/3})$. The main challenge in this context is to bound the number of distinct contributions of the type $(b_{i,j}, C)$. (Note that the second type of contribution, where $x \in \mathcal{C}$ is a cluster, is easily bounded, since there are only $O(n^{2/3})$ clusters). The replacement paths constructed earlier are designed so that for every cluster C there are at most $O(n^{2/3})$ distinct divergence point $b_{i,j}$ that can be paired with C and contribute to the value of some detour $P^-_{i,j} \in \mathcal{B}$. This establishes the sparsity of our `FT-ABFS` structure.

Let us note that in our application of the path-buying procedure, paths with a sufficiently large value to cost ratio are added *entirely* to the current spanner, and adding just the last edge of each of these paths will not suffice. This does not contradict Obs. 4.2,

for the following reason. Consider step τ and the current candidate path $P_\tau = P^-_{i,j}$ with a sufficiently large value to cost ratio with respect to H_τ. If one adds solely the last edge of P_τ to H_τ, then although the pair (i,j) would be satisfied in $H_{\tau+1} = H_\tau \cup \{\texttt{LastE}(P_\tau)\}$, this would distort the value and cost functions of the subsequent candidate paths in a way that would force us to add many new last edges to the spanner. Specifically, since only the last edge of P_τ was added to $H_{\tau+1}$, the subgraph $H_{\tau+1}$ enjoys almost none of the value of P_τ in improving the pairwise distances, and as a result, many subsequent paths would have a larger value (and hence might be suitable for purchasing) in comparison to their value when the path P_τ is added in its entirety to $H_{\tau+1}$. Fig. 12 illustrates the high level distinctions between the three constructions of additive spanners based on the path-buying technique: (a) The 6-additive construction of [4], (b) The 2-additive subsetwise construction of [10], where the stretch constraint is imposed only on a set of vertex pairs $S \times S$ for a given $S \subseteq V$, and (c) The 4-additive FT-ABFS structure presented here.

Recall that $\pi(C_k) = \pi(s, \texttt{LCA}(C_k))$ is the maximal shortest path segment shared by all the members of the cluster. For every $e_j \in T_0$, define $V_j(\mathcal{F}_3) = \{v_i \mid \texttt{LastE}(P^*_{i,j}) \notin \mathcal{F}_3\}$. By Lemma 4.16(c), it then holds that $b_{i,j} \in \pi^{mid}(C(v_i))$ for every $v_i \in V_j(\mathcal{F}_3)$. For every missing ending path $P^*_{i,j}$ (i.e., $\texttt{LastE}(P^*_{i,j}) \notin \mathcal{F}_3$) let $P^-_{i,j} = P^*_{i,j}[b_{i,j}, v_i]$ be the detour segment starting from the unique divergence point $b_{i,j}$. Let

$$\mathcal{P} = \{P^-_{i,j} \mid e_j \in T_0, \texttt{LastE}(P^*_{i,j}) \notin \mathcal{F}_3\} \tag{12}$$

be the candidate paths to be bought and added to $\mathcal{F}_3$. The benefit of considering these detours is that by the definition of $\mathcal{F}_3$, we are guaranteed that $P^-_{i,j} \subseteq G$ and in addition, each of these detours $P^-_{i,j}$ is edge disjoint with $\pi(s, \texttt{LCA}(C_i))$ for every cluster C_i that has a vertex $v' \in C_i \cap P^-_{i,j}$ with missing edge $\texttt{LastE}(P^-_{i,j}[b_{i,j}, v']) \notin \mathcal{F}_3$. The property is heavily exploited in both the size and correctness analysis of the path-buying procedure.

The Scheme. The path-buying scheme is as follows. Starting with $H_0 = \mathcal{F}_3 \setminus T_0$, the paths of $\mathcal{P} = \{P^-_{i,j}\}$ are considered in an arbitrary order. At step t, we are given $H_t \subseteq G$ and consider the $b_{i,j} - v_i$ detour $P^-_{i,j} \in G \setminus \{e_j\}$. To decide whether $P^-_{i,j}$ should be added to H_t, the cost and value of $P^-_{i,j}$ are computed as follows. Let $\texttt{Cost}(P^-_{i,j}) = |P^-_{i,j} \setminus E(H_t)|$ be the number of edges of $P^-_{i,j}$ that are missing in the current subgraph H_t. $\texttt{Cost}(P^-_{i,j})$ thus represents the increase in the size of the current FT-ABFS structure H_t if the procedure adds $P^-_{i,j}$. Let $\kappa = \texttt{Cost}(P^-_{i,j})$ and let $\widetilde{e}_1, \ldots, \widetilde{e}_\kappa$ be the edges in $P^-_{i,j} \setminus E(H_t)$ directed away from $b_{i,j}$ where $\widetilde{e}_\ell = (y_\ell, z_\ell)$.

Define $Z_{i,j} = \{z_{3\ell+1} \mid \ell \in \{0, \ldots, \lfloor(\kappa-1)/3\rfloor\}\} \subseteq P^-_{i,j}$ be the endpoints of the missing edges on $P^-_{i,j}$. Hence, it holds that

$$\texttt{LastE}(P^-_{i,j}[b_{i,j}, z_\ell]) \notin H_t \text{ for every } z_\ell \in Z_{i,j} \text{ and } \mathrm{dist}(z_\ell, z_{\ell'}, P^-_{i,j}) \;\geq\; 3$$

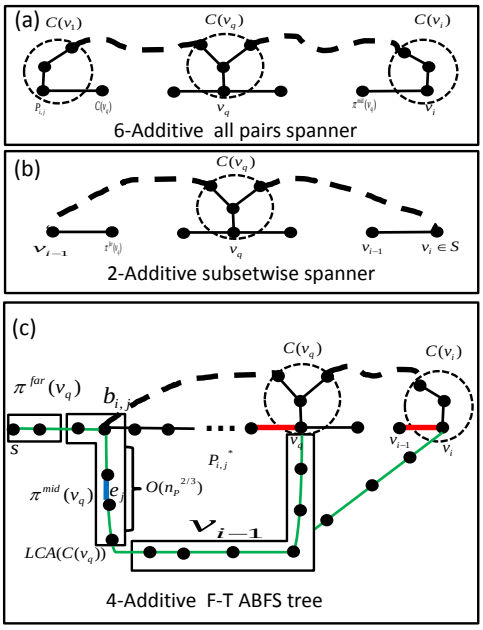

Figure 12: The green line represents the BFS edges. The edge e_j is the failed edge. Missing edges on the missing-ending path $P^*_{i,j}$ are drawn in red. Only the detour segment $P^-_{i,j} = P^*_{i,j}[b_{i,j}, v_i]$ of the path is a candidate for buying. The edges of v_i appearing on $P^*_{i,j'}$ for $e_{j'} \in \pi^{far}(v_i) \cup \pi^{near}(v_i)$ are already included in the preliminary subgraph constructed prior to the path-buying procedure. The shortest inter-cluster paths in H are indicated by dashed curves, where H is the current subgraph at the time when $P^-_{i,j}$ is considered to be bought. For each cluster $C_q \in \mathcal{C}$, with a vertex $v_q \in C(v_q)$ ending with a missing edge, the branching point (first vertex on the detour) belongs to $\pi^{mid}(v_q)$. Since the size of $\pi^{mid}(v_q)$ is bounded by $O(n^{2/3})$, the contribution of pairs $(b_{i,j}, C(v_q))$ is bounded in a similar manner to the subsetwise case of [10]. In addition, since the endpoint v_i of the detour is treated as a cluster, $C(v_i)$, the contribution of pairs $(C(v_q), C(v_i))$ is bounded by the number of clusters, $O(n^{2/3})$.

for every $z_\ell, z_{\ell'} \in Z_{i,j}$. Let

$$Cont_b(P^-_{i,j}) = \{(b_{i,j}, C_\ell) \mid \exists z_\ell \in C_\ell \cap Z_{i,j} \text{ s.t} \\ \text{dist}(b_{i,j}, z_\ell, P^-_{i,j}) < \text{dist}(b_{i,j}, C_\ell, H_t \setminus \pi(C_\ell))\} \tag{13}$$

be the set of pairs $(b_{i,j}, C_\ell)$ such that the distance between $b_{i,j}$ and C_ℓ is improved by adding $P^-_{i,j}$ to $H_t \setminus \pi(C_\ell)$. Similarly, let $C_{\ell'} = C(v_i)$ be the cluster of v_i and define

$$Cont_v(P^-_{i,j}) = \{(v_i, C_\ell) \mid \exists z_\ell \in C_\ell \cap Z_{i,j} \text{ s.t} \\ \text{dist}(v_i, z_\ell, P^-_{i,j}) < \text{dist}(C_{\ell'}, C_\ell, H_t \setminus \pi(C_{\ell'}))\} \tag{14}$$

as the set of pairs (v_i, C_ℓ) such that the distance between C_ℓ and $C_{\ell'}$ is improved by adding $P^-_{i,j}$ to H_t. Define

$$\texttt{Val}(P^-_{i,j}) = |Cont_b(P^-_{i,j})| + |Cont_v(P^-_{i,j})| \ . \tag{15}$$

The path-buying strategy is as follows. If

$$\texttt{Cost}(P^-_{i,j}) \leq 4 \cdot \texttt{Val}(P^-_{i,j}) \tag{16}$$

then we "buy" the path $P^-_{i,j}$, namely, set $H_{t+1} = H_t \cup P^-_{i,j}$. Otherwise, we do not buy $P^-_{i,j}$ and set $H_{t+1} = H_t$. The final spanner is given by

$$H = (H_{t'} \cup \mathcal{F}_3) \text{ where } t' = |\mathcal{P}| + 1.$$

4.4.1 Analysis

We begin with the following observation.

Observation 4.18 *For every $P^-_{i,j} \in \mathcal{P}$ and for every $z_\ell, z_{\ell'} \in Z_{i,j}$, it holds that $C(z_\ell) \neq C(z_{\ell'})$.*

Proof: Recall that by Fact 4.3(2), the diameter of each cluster is at most 2. As the distance between different z_ℓ's in $Z_{i,j}$ is at least 3, it follows that z_ℓ and $z_{\ell'}$ belong to distinct clusters. ∎

To establish the correctness of the construction, we show the following.

Lemma 4.19 *For every $e_j \in T_0$ and $v_i \in V(G)$ it holds that* $\text{dist}(s, v_i, H \setminus \{e_j\}) \leq \text{dist}(s, v_i, G \setminus \{e_j\}) + 4$.

Proof: The proof is by contradiction. Let $BP = \{(i,j) \mid v_i \in V(G), e_j \in \pi(s, v_i)$ and $\text{dist}(s, v_i, H \setminus \{e_j\}) > \text{dist}(s, v_i, G \setminus \{e_j\}) + 4\}$ be the set of "bad pairs," namely, vertex-edge pairs (i,j) whose additive stretch in H is greater than 4. Assume, towards contradiction, that $BP \neq \emptyset$. For each bad pair $(i,j) \in BP$, it holds that $e_j \in \pi(s, v_i)$ (since $T_0 \subseteq H$). For every bad pair $(i,j) \in BP$ define $\widetilde{P}_{i,j} \in SP(s, v_i, G \setminus \{e_j\})$ to be the replacement path whose last missing edge in H, $\widetilde{e}_{i,j}$, is the shallowest among all $s - v_i$ replacement paths in $G \setminus \{e_j\}$. Let $d(i,j) = \text{dist}(s, \widetilde{e}_{i,j}, \widetilde{P}_{i,j})$. Finally, let $(i_0, j_0) \in BP$ be the pair that minimizes $d(i,j)$, and let $\widetilde{e}_{i_0,j_0} = (v_{i'}, v_{i_1})$. Note that $\widetilde{e}_{i_0,j_0}$ is the *shallowest* "deepest missing edge" over all bad pairs $(i,j) \in BP$. See Fig. 13 for an illustration.

Claim 4.20 $(i_1, j_0) \in BP$.

Proof: Assume towards contradiction that $(i_1, j_0) \notin BP$ and let $P'' \in SP(s, v_{i_1}, H \setminus \{e_{j_0}\})$. Hence, since $(i_1, j_0) \notin BP$, it holds that

$$\begin{aligned} |P''| &\leq \text{dist}(s, v_{i_1}, G \setminus \{e_{j_0}\}) + 4 \\ &= |P^*_{i_0,j_0}[s, v_{i_1}]| + 4. \end{aligned} \tag{17}$$

We now consider the following $s - v_{i_0}$ replacement path $Q = P'' \circ \widetilde{P}_{i_0,j_0}[v_{i_1}, v_{i_0}]$. By definition of (i_1, j_0), $Q \subseteq G \setminus \{e_{j'}\}$. In addition,

$$\begin{aligned} |Q| &= |P''| + |\widetilde{P}_{i_0,j_0}(v_{i_1}, v_{i_0})| \leq |\widetilde{P}_{i_0,j_0}[s, v_{i_1}]| + 4 + |\widetilde{P}_{i_0,j_0}(v_{i_1}, v_{i_0})| \\ &= |\widetilde{P}_{i_0,j_0}| + 4 \;=\; \text{dist}(s, v_{i_0}, G \setminus \{e_{j_0}\}) + 4 \;, \end{aligned}$$

where the inequality follows by Eq. (17). This contradicts the fact that $(i_1, j_0) \in BP$. ∎

In particular, we have that $e_{j_0} \in \pi(s, v_{i_1})$, hence $\widetilde{P}_{i_1,j_0}$ is defined.

Claim 4.21 $\texttt{LastE}(\widetilde{P}_{i_1,j_0}) \notin H$.

Proof: Assume towards contradiction that $\texttt{LastE}(\widetilde{P}_{i_1,j_0}) \in H$. Then since $(i_1, j_0) \in BP$ it holds that there exists at least one missing edge in $\widetilde{P}_{i_1,j_0}$ which is strictly above $\texttt{LastE}(\widetilde{P}_{i_1,j_0})$. Since $|\widetilde{P}_{i_1,j_0}| = |\widetilde{P}_{i_0,j_0}[s, v_{i_1}]|$, we get that $d(i_1, j_0) < d(i_0, j_0)$, contradiction to the definition of (i_0, j_0). ∎

Since the last edge of $\widetilde{P}_{i_1,j_0}$ is missing, it holds that (a) v_{i_1} is clustered (by Fact 4.3) and (b) that the last edge of every $s - v_{i_1}$ replacement path is missing in H as well. In particular, for the replacement path $P^*_{i_1,j_0}$ constructed by Algorithm $\texttt{Pcons}$, $\texttt{LastE}(P^*_{i_1,j_0})$ is missing in H. Hence, it holds that $e_{j_0} \in \pi^{mid}(v_{i_1})$. It then holds by Eq. (12), that $P^-_{i_1,j_0} = P^*_{i_1,j_0}[b_{i_1,j_0}, v_{i_1}] \in \mathcal{P}$, i.e., $P^-_{i_1,j_0}$ is in the collection of detours considered to be purchased in the path-buying Procedure. Let $P' = P^-_{i_1,j_0}$ and t be the iteration where P' was considered to be added to H_t. We now consider two cases. If P' was bought, then $P^*_{i_1,j_0} = \pi(s, b_{i_1,j_0}) \circ P' \subseteq H$ hence we get a contradiction to the fact that $(i_1, j_0) \in BP$. Hence, it remains to consider the case where P' was not bought.

Let $\kappa = \texttt{Cost}(P')$ and $Z_{i_1,j_0} = \{z_1, \ldots, z_{\kappa'}\}$ for $\kappa' = \lfloor \kappa/3 \rfloor$ be the corresponding vertices in P' with a missing edges that satisfy Eq. (13). By Obs. 4.18, each $z_\ell \in Z_{i_1,j_0}$ belongs to a distinct cluster C_ℓ. Hence there are at least κ' distinct clusters on P'. Let C_r be the cluster of v_{i_1} (since v_{i_1} is incident to a missing edge in $G_C \subseteq H$, by Fact 4.3(1), v_{i_1} is indeed clustered).

A cluster C_ℓ with $z_\ell \in Z_{i_1,j_0}$ is a *contributor* if adding P' to H_t improves either the $b_{i_1,j_0} - C_\ell$ distance or the $v_{i_1} - C_\ell$ distance in H_t. I.e., if it satisfies either $\text{dist}(b_{i_1,j_0}, z_\ell, P') < \text{dist}(b_{i_1,j_0}, C_\ell, H_t \setminus \pi(C_\ell))$ (hence $(b_{i_1,j_0}, C_\ell) \in Cont_b(P')$), or $\text{dist}(v_{i_1}, z_\ell, P') < \text{dist}(C_r, C_\ell, H_t \setminus \pi(C_r))$ (hence $(b_{i_1,j_0}, C_\ell) \in Cont_v(P')$). Otherwise, C_ℓ

is *neutral*. That is, C_ℓ is neutral if $(b_{i_1,j_0}, C_\ell) \notin Cont_b(P')$ and in addition, $(v_{i_1}, C_\ell) \notin Cont_v(P')$. There are two cases to consider. If all clusters are contributors (i.e., there is no neutral cluster) then all the κ' clusters contribute to $\texttt{Val}(P')$ (either with b_{i_1,j_0} or with v_{i_1} or both). It then holds that $\texttt{Val}(P') \geq \kappa' \geq \texttt{Cost}(P')/4$. Hence, by Eq. (16), we get a contradiction to the fact that P' was not added to H_t. In the other case, there exists at least one neutral cluster C_ℓ having a unique vertex z_ℓ in P' such that

$$\begin{aligned} \text{dist}(b_{i_1,j_0}, C_\ell, H_t \setminus \pi(C_\ell)) &\leq \text{dist}(b_{i_1,j_0}, z_\ell, P') \text{ and} \quad (18) \\ \text{dist}(C_r, C_\ell, H_t \setminus \pi(C_r)) &\leq \text{dist}(v_{i_1}, z_\ell, P') \ . \quad (19) \end{aligned}$$

Let $u \in C_\ell$ be the closest vertex in the cluster C_ℓ to the divergence point b_{i_1,j_0} in the graph $H_t \setminus \pi(C_\ell)$ and define $Q_1 \in SP(b_{i_1,j_0}, u, H_t \setminus \pi(C_\ell))$ such that $u \in C_\ell$, hence $|Q_1| = \text{dist}(b_{i_1,j_0}, C_\ell, H_t \setminus \pi(C_\ell))$. Let $u' \in C_\ell$, $y' \in C_r$ be the closest pair of vertices in the cluster C_ℓ and C_r repetitively in the graph $H_t \setminus \pi(C_r)$. Define $Q_2 \in SP(u', y', H_t \setminus \pi(C_r))$, then $|Q_2| = \text{dist}(C_r, C_\ell, H_t \setminus \pi(C_r))$.

Let $\widehat{Q}_1 \in SP(u', u, G_C)$ and $\widehat{Q}_2 \in SP(y', v_{i_1}, G_C)$. Since $y', v_{i_1} \in C_r$ and $u, u' \in C_\ell$, it holds that $|\widehat{Q}_1|, |\widehat{Q}_2| \leq 2$. Consider the following $b_{i_1,j_0} - v_{i_1}$ replacement path $P_4 = Q_1 \circ \widehat{Q}_1 \circ Q_2 \circ \widehat{Q}_2$. We first claim that $P_4 \subseteq H \setminus \{e_{j_0}\}$. First note that since $G_C \subseteq (G \setminus T_0)$, it holds that the intra cluster paths $\widehat{Q}_1$ and $\widehat{Q}_2$ are free of the failed edge e_{j_0}. In addition, since $\texttt{LastE}(P^*_{i_1,j_0}), \texttt{LastE}(P^*_{i_1,j_0}[s, z_\ell]) \notin \mathcal{F}_3$, by Lemma 4.16(c), it holds that $e_{j_0} \in \pi(C_\ell) \cap \pi(C_r)$. In addition, by the definition of P_4, it holds that $P_4 \subseteq H_t \setminus (\pi(C_\ell) \cap \pi(C_r))$.

Finally, we bound the length of P_4.

$$|P_4| \ = \ |Q_1| + |Q_2| + 4 \leq \text{dist}(b_{i_1,j_0}, z_\ell, P') + \text{dist}(v_{i_1}, z_\ell, P') + 4 = |P'| + 4 \ ,$$

where the first inequality follows by Eq. (18). We therefore have that the path $P_5 = \pi(s, b_{i_1,j_0}) \circ P_4$ exists in $H \setminus \{e_{j_0}\}$ and in addition,

$$|P_5| \leq |P^*_{i_1,j_0}[s, b_{i_1,j_0}]| + |P^-_{i_1,j_0}| + 4 = |P^*_{i_1,j_0}| + 4.$$

This contradicts the assumption that $(i_1, j_0) \in BP$. The Lemma follows. ∎

Finally, we bound the size of H.

Lemma 4.22 $|E(H)| \leq n^{4/3}$.

Proof: Let $\mathcal{B} \subseteq \mathcal{P}$ be the set of paths that were bought in the path-buying procedure. Since $|E(\mathcal{F}_3) \setminus T_0| = O(n^{4/3})$, it remains to bound the number of edges added due to the paths in $\mathcal{B}$. Note that for $\gamma = 1/3$, it holds by Fact 4.3, that there are $O(n^{2/3})$ clusters. In addition, recall that for every cluster C_i, it holds that $|\pi^{mid}(C_i)| = O(n^{2/3})$. We have

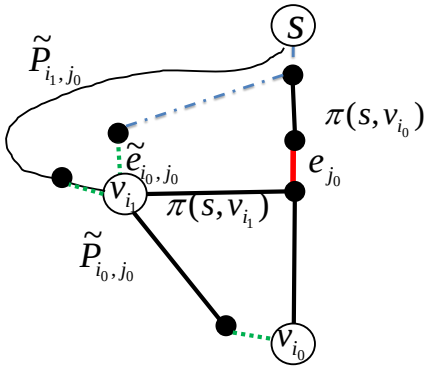

Figure 13: The dashed edges are missing in H. Since (i_0, j_0) is a bad pair, it holds that (i_1, j_0) is a bad pair as well and in addition, the last edge of $\widetilde{P}_{i_1,j_0}$ is missing in H.

the following.

$$\begin{aligned} |H \setminus \mathcal{F}_3| &= \sum_{P^-_{i,j} \in \mathcal{B}} \texttt{Cost}(P^-_{i,j}) \leq \sum_{P^-_{i,j} \in \mathcal{B}} 4 \cdot \texttt{Val}(P^-_{i,j}) \\ &= 4 \sum_{P^-_{i,j} \in \mathcal{B}} \left(|Cont_b(P^-_{i,j})| + |Cont_v(P^-_{i,j})| \right) , \end{aligned} \tag{20}$$

where the inequality follows by Eq. (16), and the last equality follows by Eq. (15).

Our counting strategy is as follows. We fix a cluster C_ℓ, and bound the number of times it appears in either $Cont_b(P')$ or $Cont_v(P')$ over all $P' \in \mathcal{B}$. The following definition is useful in our analysis. Define the multisets

$$\begin{aligned} Cont_b(C_\ell) &= \{(b_{i,j}, C_\ell) \mid P^-_{i,j} \in \mathcal{B} \text{ and } (b_{i,j}, C_\ell) \in Cont_b(P^-_{i,j})\}, \\ Cont_v(C_\ell) &= \{(v_i, C_\ell) \mid \exists P^-_{i,j} \in \mathcal{B} \text{ and } (v_i, C_\ell) \in Cont_v(P^-_{i,j})\}. \end{aligned}$$

Note that certain pairs might appear several times in both $Cont_b(C_\ell)$ and $Cont_v(C_\ell)$. Hence, we do not bound the number of unique pairs but also take into account the possible reappearance of the same pair in these expressions. Eq. (20) can now be written equivalently as

$$|H \setminus \mathcal{F}_3| \leq 4 \sum_{P^-_{i,j} \in \mathcal{B}} \left(|Cont_b(P^-_{i,j})| + |Cont_v(P^-_{i,j})| \right) = 4 \sum_{C_\ell} |Cont_b(C_\ell)| + |Cont_v(C_\ell)| .$$

We now show the following.

Lemma 4.23 *$|Cont_b(C_\ell)| = O(n^{2/3})$ for every C_ℓ.*

Proof: We begin by showing that for $P^-_{i,j} \in \mathcal{B}$, if $(b_{i,j}, C_\ell) \in Cont_b(P^-_{i,j})$ then $b_{i,j} \in \pi^{mid}(C_\ell)$. To see that, let t be the iteration in which $P^-_{i,j}$ was added to H_t. It then holds that there exists $z_\ell \in C_\ell \cap Z_{i,j}$ that satisfies Eq. (13). By the definition of $Z_{i,j}$, we have that $\texttt{LastE}(P^*_{i,j}[s, z_\ell]) \notin \mathcal{F}_3$. Hence, by Lemma 4.16(c) it holds that $b_{i,j} \in \pi^{mid}(C_\ell)$.

We proceed by showing that for every fixed $b_{i,j} \in \pi^{mid}(C_\ell)$, the pair $(b_{i,j}, C_\ell)$ can appear at most 3 times in the multiset $Cont_b(C_\ell)$. Formally, define the set of paths in which $b_{i,j}$ contributes with the cluster C_ℓ as

$$A(i,j) \;=\; \{P^-_{i',j'} \in \mathcal{B} \;\mid\; b_{i',j'} = b_{i,j} \;\text{ and }\; (b_{i',j'}, C_\ell) \in Cont_b(P^-_{i',j'})\}\ .$$

Claim 4.24 *$|A(i,j)| \leq 3$ for every i, j such that $b_{i,j} \in \pi^{mid}(C_\ell)$.*

Proof: Let $A(i,j) = \{Q_1 = P^-_{i_1,j_1}, \ldots, Q_N = P^-_{i_N,j_N}\}$ be sorted according to the time t_k they were bought and added to the graph H_{t_k}. Hence, $b_{i_1,j_1} = \ldots = b_{i_N,j_N} = b_{i,j}$. We claim that $N \leq 3$. For each $k \in \{1, \ldots, N\}$, considered at time t_k where $t_k < t_{k+1}$, there exists a vertex $z_k \in Z_{i_k,j_k} \cap C_\ell$ such that $\texttt{LastE}(Q_k[b_{i_k,j_k}, z_k]) \notin H_{t_k}$ and $\text{dist}(b_{i,j}, z_k, Q_k) = \text{dist}(b_{i_k,j_k}, z_k, Q_k) < \text{dist}(b_{i,j}, C_\ell, H_{t_k} \setminus \pi(C_\ell))$. This holds as by definition, $b_{i,j} = b_{i_k,j_k}$ for every $k \in \{1, \ldots, N\}$ and by Eq. (13). We now show that, denoting

$$Y_k = \text{dist}(b_{i,j}, z_k, H_{t_{k+1}} \setminus \pi(C_\ell)),$$

we have $Y_k < Y_{k-1}$ for every $k \in \{2, \ldots, N\}$.

At each time t_k, since a contribution is made, and by the fact that $z_k, z_{k-1} \in C_\ell$, we have that

$$\begin{aligned} Y_k \;&\leq\; \text{dist}(b_{i,j}, z_k, Q_k \setminus \pi(C_\ell)) = \text{dist}(b_{i,j}, z_k, Q_k) \\ &<\; \text{dist}(b_{i,j}, C_\ell, H_{t_k} \setminus \pi(C_\ell)) \leq Y_{k-1}\ , \end{aligned} \tag{21}$$

where the first inequality holds as Q_k was bought at time t_k and hence $Q_k \subseteq H_{t_{k+1}}$. By Obs. 4.18, $z_{k-1}, z_k \in C_\ell$ are the unique vertices in C_ℓ in the sets $Z_{i_{k-1},j_{k-1}}$ and Z_{i_k,j_k} respectively. The equality of (21) holds since $\texttt{LastE}(P^*_{i_k,j_k}[s, z_k]) \notin \mathcal{F}_3$ and hence by 4.16(b), Q_k is edge disjoint with $\pi(s, z_k)$. Since by Obs. 4.4, $\pi(C_\ell) \subseteq \pi(s, z_k)$, it holds that Q_k and $\pi(C_\ell)$ are edge disjoint as well. The strict inequality in (21) follows by Eq. (13) and by the fact that the pair (z_k, C_ℓ) contributes to the value of Q_k i.e., $(b_{i,j}, C_\ell) \in Cont_b(Q_k)$. The last inequality follows by the fact that $z_{k-1} \in C_\ell$. Hence, letting

$$Y = \text{dist}(b_{i,j}, z_N, H_{t_{N+1}} \setminus \pi(C_\ell)),$$

we have

$$Y \leq \text{dist}(b_{i,j}, z_1, H_{t_2} \setminus \pi(C_\ell)) - (N-1). \tag{22}$$

Conversely, we also have the following.

$$\begin{aligned} Y &\geq \text{dist}(b_{i,j}, z_N, H_{t_{N+1}} \setminus \{e_{j_1}\}) && (23)\\ &\geq \text{dist}(b_{i,j}, z_1, H_{t_{N+1}} \setminus \{e_{j_1}\}) - 2 && (24)\\ &\geq \text{dist}(b_{i,j}, z_1, G \setminus \{e_{j_1}\}) - 2 && (25)\\ &= \text{dist}(b_{i,j}, z_1, Q_1) - 2 && (26)\\ &= \text{dist}(b_{i,j}, z_1, Q_1 \setminus \pi(C_\ell)) - 2 && (27)\\ &\geq \text{dist}(b_{i,j}, z_1, H_{t_2} \setminus \pi(C_\ell)) - 2\ , && (28) \end{aligned}$$

where Eq. (23) follows by the fact that $z_1 \in C_\ell$, $\texttt{LastE}(P^*_{i_1,j_1}[s, z_1]) \notin \mathcal{F}_3$ and hence by Lemma 4.16(a), $e_{j_1} \in \pi(C_\ell)$. Eq. (24) follows by the fact that $z_1, z_N \in C_\ell$ and by Fact 4.3(2). Eq. (25) follows by the fact that $H_{t_{N+1}} \subseteq G$. Eq. (26) follows by the fact that $P^*_{i_1,j_1} \in SP(s, v_{i_1}, G \setminus \{e_{j_1}\})$ and $Q_1[b_{i,j}, z_1] = P^*_{i_1,j_1}[b_{i,j}, z_1]$. To see Eq. (27), note that by Lemma 4.16(b), it holds that $Q_1[b_{i,j}, z_1]$ and $\pi(s, z_1)$ are edge disjoint. In addition, by Obs. 4.4, $\pi(C_\ell) \subseteq \pi(s, z_1)$, hence $Q_1[b_{i,j}, z_1]$ and $\pi(C_\ell) \subseteq \pi(s, z_1)$ are edge disjoint as well. Finally, Eq. (28) follows as $Q_1 \subseteq H_{t_2}$ since $Q_1 \in \mathcal{B}$ was bought at time t_1.

We therefore have that $\text{dist}(b_{i,j}, z_N, H_{t_{N+1}} \setminus \pi(C_\ell)) \geq \text{dist}(b_{i,j}, z_1, H_{t_2} \setminus \pi(C_\ell)) - 2$, combining with Eq. (22), we have $|A(i,j)| = N \leq 3$. ∎

Since $A(i,j) \neq \emptyset$ only if $b_{i,j} \in \pi^{mid}(C_\ell)$, it follows from Cl. 4.24 that

$$|Cont_b(C_\ell)| \ = \sum_{b_{i,j} \in \pi(s,v_i)} |A(i,j)| \ = \sum_{b_{i,j} \in \pi^{mid}(C_\ell)} |A(i,j)| = O(n^{2/3}).$$

The lemma follows. ∎

We now turn to consider the second type of contribution of the form (v_i, C_ℓ).

Lemma 4.25 *$|Cont_v(C_\ell)| = O(n^{2/3})$ for every C_ℓ.*

Proof: For every cluster $C_r \neq C_\ell$ define the multiset

$$D_r = \{(v_i, C_\ell) \in Cont_v(C_\ell) \mid v_i \in C_r\}.$$

Note that the pair (v_i, C_ℓ) can contribute several times to D_r with the same v_i. Hence, we do not count only unique pairs as the same pair might contribute several times. In fact, we count the number of all times in which a path $P' \in \mathcal{B}$ was bought, and the pair $(v_i, C_\ell) \in Cont_v(P')$ such that $v_i \in C_r$ contributes to the value of the path P'. Since missing edges in G_C are between vertices of different clusters, we have that $Cont_v(C_\ell) = \bigcup_{r \neq \ell} D_r$. We now show that each $|D_r| \leq 5$ which concludes the proof since there are overall $|\mathcal{C}| = O(n^{2/3})$ clusters in G_C.

Let $Q_1 = P^-_{i_1,j_1}, \ldots, Q_N = P^-_{i_N,j_N}$ be such that $(v_{i_k}, C_\ell) \in D_r$, where the paths are sorted according to the time t_k they were considered, for every $k \in \{1, \ldots, N\}$ and $N = |D_r|$.

We do not assume that the $v_{i,k'}$'s are distinct. Let $z_k \in Z_{i_k,j_k} \cap C_\ell$ be the unique vertex of C_ℓ in Z_{i_k,j_k} that contributes by adding its path Q_k to v_{i_k} to H_{t_k}. We then have that $\texttt{LastE}(Q_k[b_{i_k,j_k}, z_k]) \notin H_{t_k}$ and hence by Lemma 4.16(a), $e_{j_k} \in \pi(C_\ell)$. In addition, since $\texttt{LastE}(Q_k) \notin \mathcal{F}_3$, by Lemma 4.16(a) it also holds that $e_{j_k} \cap \pi(C_r)$ as $v_{i_k} \in C_r$. Hence,

$$e_{j_k} \in \pi(C_\ell) \cap \pi(C_r) \text{ for every } k \in \{1, \ldots, N\} \ . \tag{29}$$

We next show that, denoting

$$X_k = \text{dist}(v_{i_k}, z_k, H_{t_{k+1}} \setminus \pi(C_r)),$$

we have $X_k < X_{k-1}$ for every $k \in \{2, \ldots N\}$. Note the each v_{i_k} belongs to the same cluster C_r and every z_k for $k \in \{1, \ldots, N\}$ belongs to the same cluster C_ℓ. Each time a contribution is made at time t_k, it implies that

$$\begin{aligned}
X_k &\leq \text{dist}(v_{i_k}, z_k, Q_k \setminus \pi(C_r)) && (30)\\
&= \text{dist}(v_{i_k}, z_k, Q_k) && (31)\\
&< \text{dist}(C_r, C_\ell, H_{t_k} \setminus \pi(C_r)) && (32)\\
&\leq X_{k-1} \ , && (33)
\end{aligned}$$

where Eq. (30) follows by the fact that $Q_k \subseteq H_{t_{k+1}}$, Eq. (31) follows by the fact that $\texttt{LastE}(P^*_{i_k,j_k}) \notin \mathcal{F}_3$, and hence by Lemma 4.16, $Q_k = P^*_{i_k,j_k}[b_{i_k,j_k}, v_{i_k}]$ and $\pi(s, v_{i_k})$ are edge disjoint. In addition, by Obs. 4.4, $\pi(C_r) \subseteq \pi(s, v_{i_k})$, hence $\pi(C_r)$ and Q_k are edge disjoint as well. Eq. (32) follows by Eq. (14) and the fact that $(v_{i_k}, C_\ell) \in Cont_v(Q_k)$ and finally, Eq. (33) follows by the fact that $v_{i_{k-1}} \in C_r$ and $z_{k-1} \in C_\ell$. Therefore, letting

$$X = \text{dist}(v_{i_N}, z_N, H_{t_{N+1}} \setminus \pi(C_r)),$$

we have that

$$X \leq \text{dist}(v_{i_1}, z_1, H_{t_2} \setminus \pi(C_r)) - (N-1) \ . \tag{34}$$

Conversely, we have that

$$\begin{aligned}
X &\geq \text{dist}(v_{i_N}, z_N, G \setminus \pi(C_r)) && (35)\\
&\geq \text{dist}(v_{i_1}, z_1, G \setminus \pi(C_r)) - 4 && (36)\\
&= \text{dist}(v_{i_1}, z_1, Q_1) - 4 = \text{dist}(v_{i_1}, z_1, Q_1 \setminus \pi(C_r)) - 4 && (37)\\
&\geq \text{dist}(v_{i_1}, z_1, H_{t_2} \setminus \pi(C_r)\}) - 4 \ , && (38)
\end{aligned}$$

where Eq. (35) follows as $H_{t_{N+1}} \subseteq G$. Eq. (36) follows by the fact that $v_{i_1}, v_{i_N} \in C_r$ and $z_1, z_N \in C_\ell$ and by Fact 4.3(b). Eq. (37) follows by the fact that $P^*_{i_1,j_1}[z_1, v_{i_1}] = Q_1[z_1, v_{i_1}]$ and $P^*_{i_1,j_1} \in SP(s, v_{i_1}, G \setminus \{e_{j_1}\})$. In addition, by Lemma 4.16, since $\texttt{LastE}(P^*_{i_1,j_1}) \notin \mathcal{F}_3$, it holds that Q_1 and $\pi(s, v_{i_1})$ are edge disjoint. By Obs. 4.4, $\pi(C_r) \subseteq \pi(s, v_{i_1})$, hence Q_1

and $\pi(C_r)$ are edge disjoint as well. Finally, Eq. (38), follows as $Q_1 \subseteq H_{t_2}$ since $Q_1 \in \mathcal{B}$ was added to H_{t_1} at time t_1.

Combining with Eq. (34), we get that $N \leq 5$ and hence $|Cont_v(C_\ell)| = O(n^{2/3})$. The claim follows. ∎

Overall, by Lemma 4.23 and Lemma 4.25, we have that $|Cont_b(C_\ell)| + |Cont_v(C_\ell)| = O(n^{2/3})$ for every C_ℓ. Hence, by plugging into Eq. (21), as there are $O(n^{2/3})$ clusters, we have that $|H| \leq 4 \sum_{C_\ell} (|Cont_b(C_\ell)| + |Cont_v(C_\ell)|) = O(n^{4/3})$ as required. The Lemma follows. ∎

References

[1] D. Aingworth, C. Chekuri, P. Indyk, and R. Motwani. Fast estimation of diameter and shortest paths (without matrix multiplication). *SIAM J. Comput.*, 28(4):1167–1181, 1999.

[2] S. Baswana and S. Sen. Approximate distance oracles for unweighted graphs in expected $O(n^2)$ time. *ACM Trans. Algorithms*, 2(4):557–577, 2006.

[3] S. Baswana and N. Khanna. Approximate Shortest Paths Avoiding a Failed Vertex: Optimal Size Data Structures for Unweighted Graph. In *Proc. 27th Symp. on Theoret. Aspects of Computer Sci.*, 513–524, 2010.

[4] S. Baswana, T. Kavitha, K. Mehlhorn, and S. Pettie. Additive spanners and (α, β)-spanners. *ACM Trans. Algo.* 7, A.5, 2010.

[5] A. Bernstein. A nearly optimal algorithm for approximating replacement paths and k shortest simple paths in general graphs. In *Proc. 21st ACM-SIAM Symp. on Discrete Algorithms*, 2010.

[6] G. Braunschvig, S. Chechik and D. Peleg. Fault tolerant additive spanners. In *Proc. 38th Workshop on Graph-Theoretic Concepts in Computer Science*, 206–214, 2012.

[7] S. Chechik, M. Langberg, D. Peleg, and L. Roditty. f-sensitivity distance oracles and routing schemes. *Algorithmica* **63**, (2012), 861–882.

[8] S. Chechik, M. Langberg, D. Peleg, and L. Roditty. Fault-tolerant spanners for general graphs. In *Proc. 41st ACM Symp. on Theory of Computing*, 435–444, 2009.

[9] S. Chechik. New Additive Spanners. In *Proc. 24th Symp. on Discrete Algorithms*, 2013.

[10] M. Cygan, F. Grandoni and T. Kavitha. On Pairwise Spanners. In *Proc. 30th Symp. on Theoret. Aspects of Computer Sci.*, 209–220, 2013.

[11] M. Dinitz and R. Krauthgamer. Fault-tolerant spanners: better and simpler. In *Proc. ACM Symp. on Principles of Distributed Computing*, 2011, 169-178.

[12] P. Erdos. *Extremal problems in graph theory.* Theory of Graphs and its Applications (Proc. Symp. Smolenice, 1963), page 2936, 1963.

[13] C. Godsil and G. Royle. *Algebraic Graph Theory.* Springer, 2001.

[14] K. E. Mellinger and D. Mubayi. Constructions of bipartite graphs from finite geometries. *J. Graph Theory*, (2005), 1–10.

[15] M. Parter and D. Peleg. Sparse Fault-tolerant BFS trees. In *Proc. Europ. Symp. on Algorithms*, 2013.

[16] M. Pătraşcu and L. Roditty. Distance oracles beyond the Thorup-Zwick bound. In *Proc. IEEE Symp. on Foundations of Computer Science*, 815–823, 2010.

[17] M. Smid. A probabilistic construction of a dense bipartite graph with high girth. In *Unpublished Note*, 2006.

[18] R. Wenger. Extremal graphs with no C^4's, C^6's, or C^{10}'s. *J. Combin. Th., Ser. B*, 52:113–116, 1991.

[19] D.P Woodruff. Lower bounds for additive spanners, emulators, and more. In *Proc. 47th Symp. on Foundations of Computer Science*, 389–398, 2006.